Non-invasive monitoring of peripheral perfusion

An exploration of non-invasive peripheral perfusion monitoring for applications in research & healthcare

Laurence Blaxter

A thesis presented to the University of Nottingham for the degree of
Doctor of Philosophy



Department of Electrical and Electronic engineering
University of Nottingham
United Kingdom

December 2014

Non-invasive monitoring of peripheral perfusion

An exploration of non-invasive peripheral perfusion monitoring for applications in research & healthcare

Laurence Blaxter

Abstract

The need for real-time, prognostic perfusion data has been widely recognised in clinical practice and academia. Perfusion, the volumetric blood flow to tissue, is an important tool for haemodynamic monitoring, since perfusion adequacy is vital to tissue health. Perfusion can serve as both a diagnostic indicator and a fundamental research tool. However, lack of a low cost, easily applied, and non-invasive sensor technology capable of measuring peripheral perfusion in absolute volumetric units is clearly a major factor preventing the widespread clinical use of perfusion monitoring.

This thesis explores and evaluates potential techniques for the non-invasive measurement of peripheral perfusion in absolute volumetric units. Following a review of applicable technologies, a small selection of complementary optical methods are selected for study based on economic and technology readiness arguments, and a series of volunteer studies conducted to evaluate these using both off the shelf and prototype apparatus.

A study of healthy adult and child volunteers demonstrates that Laser Doppler flowmetry (LD) outperforms both pulse oximeter derived "Perfusion Index (PI)", and a novel Capillary Refill Time (CRT) measurement device by a considerable margin in a cold challenge test. The PI metric had the lowest correlation with temperature, suggesting it is the least effective perfusion analogue of the three. A novel filtering algorithm is presented for removal of artefacts from raw LD flowmetry data. CRT was found to significantly correlate with short period oscillations in LD perfusion.

A second study combined CRT with NIRS, allowing tissue haemoglobin concentrations to be estimated. Studies of the palm of healthy volunteers' hands using this device, in combination with cooling and/or venous and arterial occlusion, demonstrate both that capillary refill measures blood mobility within the capillary bed, and that NIRS with venous occlusion is an effective method for measurement of absolute perfusion.

Although LD was the most effective existing solution, and CRT has potential, only thermal techniques and NIRS with venous occlusion allow practical, low cost quantification of absolute perfusion. Development of a thermal diffusion sensor is recommended.

Acknowledgements

I would like to thank clinical research fellows Drs Sarah Hill and Caroline Henry from the school of medicine for their assistance with volunteer experiments. Thankyou also to Dr David Morris for assistance with sensor construction and to the Nottingham University Hospitals' clinical engineering team for guidance with regulatory and device safety matters.

Contents

1 Introduction **1**

1.1 Motivation 1

1.2 Existing techniques 2

1.3 Thesis structure 3

2 Thermal perfusion measurement **5**

2.1 Introduction 5

2.1.1 Concepts 5

2.1.2 Comparison with competing techniques 6

2.1.3 Current developments in thermal perfusion monitoring 7

2.2 Heat transport in biological tissue 9

2.2.1 Pennes bioheat equation 9

2.2.2 Wulff continuum model 10

2.2.3 Chen-Holmes continuum model 11

2.2.4 Weinbaum, Jiji and Lemons bioheat model 12

2.2.5 Simplified Weinbaum-Jiji model 13

2.3 Numerical modelling of the bioheat equation 14

2.3.1 Choosing a model 14

2.3.2 Simulation of a heated disc sensor 16

2.3.3 Designing a more useful sensor 21

2.4 Passive thermal flux sensing 25

2.5 Experiments in non contact thermal sensing 27

2.6 Conclusion 30

3 Use of Nuclear Magnetic Resonance for perfusion measurement 32

3.1 Introduction . . . 32

3.1.1 Nuclear Magnetic Resonance and blood flow . . . 32

3.2 Instrument design . . . 33

3.3 Pulse sequences . . . 34

3.4 RF coil driver requirements . . . 37

3.5 Signal to noise ratio . . . 38

3.6 High Homogeneity magnet configuration . . . 40

3.7 Conclusion . . . 42

4 Trans-cutaneous blood gas monitoring for perfusion sensing 44

4.1 Introduction . . . 44

4.1.1 Current clinical practice . . . 45

4.2 Theory of operation . . . 46

4.2.1 The Clark electrode . . . 46

4.2.2 Trans-cutaneous CO_2 sensors . . . 48

4.2.3 Use of the Clark electrode for H_2 sensing . . . 48

4.3 Trans-cutaneous blood gas and perfusion . . . 49

4.3.1 H_2 washout techniques . . . 49

4.3.2 Trans-cutaneous CO_2 as a perfusion indicator . . . 51

4.3.3 Trans-cutaneous O_2 as a perfusion indicator . . . 52

4.4 Conclusion . . . 54

5 Use of medical ultrasound for perfusion sensing 57

5.1 Perfusion sensing via ultrasonic imaging . . . 57

5.1.1 Current clinical techniques in medical ultrasound . . . 57

5.1.2 Prior studies of ultrasonic perfusion monitoring . . . 58

5.2 Ultrasonically modulated optical tomography (USMOT) . . . 60

5.2.1 An introduction to USMOT . . . 60

5.2.2 Potential uses in perfusion measurement . . . 60

5.3 Photoacoustic effects 61
5.3.1 Photoacoustic effect in tissue 61
5.3.2 Use of photoacoustic effects for perfusion measurement 61
5.4 Conclusion 62

6 Optical perfusion sensing **63**
6.1 Introduction 63
6.2 Laser Doppler 63
6.2.1 Reflection mode laser Doppler/Quasi-elastic light scattering (QELS) 63
6.2.2 Diffusing wave spectroscopy and Diffuse correlation spectroscopy (DWS/DCS) 65
6.2.3 Doppler coherence techniques 66
6.3 Near-Infrared spectroscopy (NIRS) 68
6.3.1 Continuous wave 68
6.3.2 Modulated 70
6.3.3 Time resolved 70
6.3.4 Coherence and optical frequency resolved 71
6.4 Use of venous and arterial occlusion with NIRS 73
6.5 Direct imaging of capillary flow (Capillaroscopy) 74
6.5.1 Orthogonal polarization imaging (OPS) 74
6.5.2 Dark field microscopy in skin 74
6.5.3 Optical coherence tomography (OCT) techniques 75
6.6 Photo-plethysmography (PPG) 75
6.6.1 PPG amplitude and perfusion 76
6.6.2 Extracting venous saturation 76
6.6.3 Multi-wavelength PPG and tissue oxygenation 78
6.7 Capillary refill 87
6.7.1 Existing clinical use 88
6.7.2 Hypotheses concerning the origin of the refill signal 89
6.8 Conclusion 91

7 Choosing and developing a candidate technique 93

7.1 Analytic Hierarchy Process (AHP) based analysis 94

7.2 Development of a continuous wave NIRS system 102

7.2.1 Front end electronic design - dual op-amp system 108

7.2.2 Front end electronic design - single op-amp system 112

7.2.3 DSP design 116

7.2.4 NIRS sensor development 122

8 Automation of the Capillary refill test with NIRS and pressure application 129

8.1 Development of an automated pressure application system 129

8.1.1 Solenoid based capillary refill system 130

8.1.2 Linear actuator based capillary refill system 135

8.2 Analysis of NIRS data from automated capillary refill 142

8.2.1 Pulsatile modulation depth fitting 143

8.2.2 Fitting a diffusion model 152

8.2.3 Fitting a differential path length model 167

9 Development of a portable NIRS data logger with a pneumatic supply for controlled pressure application 177

9.1 Development of a portable continuous wave NIRS system 178

9.1.1 Prototype data logger electronic design 180

9.1.2 Design of the data logger pneumatics 184

9.1.3 Data logger firmware development 187

9.2 Development of a pneumatic capillary refill sensor 190

9.2.1 Prototype pneumatic sensor 191

9.2.2 Development of a production ready pneumatic sensor 195

9.3 Continuous wave NIRS data logger design for low volume production . . 204

9.3.1 Electronic design 204

9.3.2 DSP bandpass filtering 207

9.3.3 Cost analysis 208

9.3.4 Enclosure design and assembly 208

10 Experimental optical perfusion measurement in healthy volunteers 211

10.1 Experimental protocols 212

10.1.1 Proposed protocol evaluating capillary refill, venous occlusion, and Masimo PI 212

10.1.2 Use of data from "PaedCEW" research project 214

10.2 Venous occlusion and NIRS 217

10.2.1 Error analysis of venous occlusion data 222

10.2.2 Analysis of the healthy volunteer data 231

10.3 Capillary refill and NIRS 232

10.4 Masimo pulse oximeter perfusion index 234

10.4.1 PI baseline and breath hold experiments 235

10.4.2 PI motion experiments 240

10.5 Results from the PaedCEW study 241

10.5.1 Moor Instruments MoorVMS laser Doppler data 241

10.5.2 Perfusion Index from Masimo Inc. Radical 7 Pulse oximeter . . . 254

10.5.3 Threshold based capillary refill time with a pneumatic PPG sensor 259

10.5.4 Summary of PaedCEW study data 279

10.6 Experimental characterisation of venous occlusion versus capillary refill . 281

10.6.1 Proof of concept experiments on a healthy volunteer 282

10.6.2 Discussion of the NIRS data during capillary refill 286

10.6.3 Experimental characterisation of capillary refill versus venous and arterial occlusion with 10°C water bath cooling 289

11 Conclusions and potential future research 310

11.1 Summary and conclusions 310

11.2 Suggested future work 311

A Further details of capillary refill study hardware 343

A.1 Pneumatic capillary refill sensor used in the PaedCEW study 344

A.2 PaedCEW experimental setup 344

B Further details of polyurethane capillary refill sensor 345

Commonly used acronyms

LD	Laser Doppler
PI	Perfusion Index
CRT	Capillary Refill Time
PU	Perfusion Units
NIRS	Near Infrared Spectroscopy
NMR	Nuclear Magnetic Resonance
USMOT	Ultrasonically Modulated Optical Tomography
QELS	Quasi-Elastic Light Scattering
DWS	Diffusing Wave Spectroscopy
DCS	Diffuse Correlation Spectroscopy
OPS	Orthogonal Polarization Imaging/Spectroscopy
OCT	Optical Coherence Tomography
PPG	Photo-PlethysmoGraphy
MAP	Mean Arterial Pressure
MVP	Mean Venous Pressure
PVR	Peripheral Vascular Resistance
AHP	Analytic Hierarchy Process
ANP	Analytic Network Process
BORC	Benefits Opportunities Risks Costs
DSP	Digital Signal Processing
MRI	Magnetic Resonance Imaging
CT	Computer Tomographic
ASL	Arterial Spin Labelling
NICE	National Institute for heath and Care Excellence
CPAP	Continuous Positive Airway Pressure
TPSF	Temporal Point Spread Function
DPF	Differential Pathlength Function
LED	Light Emitting Diode
SPAD	Single Photon Avalanche Diode
CMOS	Complimentary Metal Oxide Semiconductor
SLD	Super-luminescent Laser Diode
ECDL	External Cavity Diode Laser
RGB	Red Green Blue
TRL	Technology Readiness Level
MRL	Manufacturing Readiness Level
PID	Proportional Integral Derivative
PCB	Printed Circuit Board
BGA	Ball Grid Array
USB	Universal Serial Bus
ADC	Analogue to Digital Converter
ENOB	Effective Number of Bits

LSB	Least Significant Bit
RTOS	Real Time Operating System
DMA	Direct Memory Access
FIFO	First In First Out
FPU	Floating Point Unit
PSSR	Power Supply Rejection Ratio
PWM	Pulse Width Modulation
RMS	Root Mean Squared
SNR	Signal to Noise Ratio
MAC	Multiply ACcumulate
FIR	Finite Impulse Response
DFT	Discrete Fourier Transform
FFT	Fast Fourier Transform
QAM	Quadrature Amplitude Modulation
OFDM	Orthogonal Frequency Division Multiplexing
SWD	Serial Wire Debug
MOSFET	Metal Oxide Semiconductor Field Effect Transistor
JFET	Junction gate Field Effect Transistor
IGMF	Infinite Gain Multiple Feedback
LDO	Low DropOut
SEPIC	Single Ended Primary Inductance Converter
GPIO	General Purpose Input/Output
PSI	Pounds per Square Inch
RF	Radio Frequency
CNC	Computer Numeric Control
PET	PolyEthylene Terephthalate
PVC	PolyVinyl-Chloride
PDMS	PolyDiMethylSiloxane
PTFE	PolyTetraFluoroEthylene
RTV	Room Temperature Vulcanising
ABS	Acrylonitrile Butadiene Styrene
COTS	Commercial Off-The-Shelf
MOPP	Means Of Patient Protection
MOOP	Means Of Operator Protection
EMC	ElectroMagnetic Compatibility
EMI	ElectroMagnetic Interference
BLUE	Best Linear Unbiased Estimator
DW	Durbin-Watson
DFA	Detrended Fluctuation Analysis
PDF	Probability Density Function
GLS	Generalized Least Squares

Chapter 1

Introduction

1.1 Motivation

A practical sensor solution that allowed continuous, non-invasive perfusion monitoring to become commonplace would have clear healthcare benefits. Alongside the obvious applications in haemodynamic monitoring and assessment of perfusion adequacy, there is evidence that peripheral perfusion is indicative of overall circulatory status. For example, in circulatory failure, blood flow is diverted to protect vital organs such as the heart and brain from hypoperfusion, meaning that flow to the skin, subcutaneous tissue, muscle, and gastrointestinal tract can decrease significantly. In a 2009 study of 50 intensive care unit admissions, Lima et al found that peripheral perfusion was a significant prognostic indicator for patients with circulatory shock [1]. In a 1996 study of postoperative care, Boyd and Bennett state that "it may no longer be acceptable to undertake surgery in these patients (those with poor cardiovascular reserve) without facilities to monitor and improve cardiac output and tissue perfusion." [2]. Shoemaker et al reached similar conclusions in a 1999 study comparing outcome in high-risk elective surgery patients [3] "the continuous on-line real-time displays of haemodynamic trends facilitate early recognition of acute circulatory dysfunction."

The importance of perfusion adequacy assessment in burn injuries or skin grafts is widely recognised, for example Furnas discusses the application of various perfusion monitoring methods to transferred free tissue [4], and Abel and Mitchell state that "If the compromised circulation of a free flap cannot be re-established within 8-12 h, salvage of such free-tissue transfer may be impossible because of the "no-reflow" phenomenon" [5].

It is thus clear that peripheral perfusion monitoring could be applied to a broad range

of applications; as a research tool in haemodynamic studies, as an assessment tool in healthcare, and as a component part of a "track and trigger" system to identify risk of vital tissue hypoperfusion due to circulatory dysfunction.

1.2 Existing techniques

Currently, numerous research level techniques are in use for the assessment of perfusion, but the majority of these are invasive, involving injection of blood flow tracers, such as fluorescent or radioactive agents, or micro-bubbles that can be traced using ultrasound.

Arterial spin labelled MRI, (ASL): allows visualisation of volumetric blood flow in three dimensions, making it an extremely powerful technique. However, very large and expensive MRI hardware is required, meaning ASL is not a widely used clinically, and is wholly unsuitable for continuous monitoring.

Laser Doppler velocimetry: an established technology, but has seen limited clinical use, possibly because it is only capable of measuring changes in perfusion in the upper capillary bed as a metric expressed in arbitrary units. It may be of use in trend measurement, and is recommended by NICE (National Institute for Health and care Excellence) for use in burn depth assessment [6].

Pulse oximetry devices marketed by Masimo Inc: feature a "Perfusion Index" output in arbitrary units, that is claimed to be of use in perfusion trend analysis.

Doppler ultrasound: devices capable of measuring blood flow velocity are available off-the-shelf, but current hardware only allows accurate perfusion measurement in highly perfused organs (e.g. liver). Off-the-shelf hardware commonly features signal processing capable of selecting a single large artery or vein in the Doppler image for the calculation of total volumetric flow versus time. However, due to the limited resolution of the raw Doppler image, this technique is only suitable for the assessment of flow through the larger vessels, for example whole limb perfusion.

Thermal perfusion measurement: these techniques have several clear advantages; the heat capacity of blood is very high, making even a small flow of blood thermally significant, thermal sensors are low cost and readily available, and although the thermal diffusion coefficient of tissue is typically over an order of magnitude greater than the diffusion coefficient of oxygen, similar transport mechanisms are involved in thermal and oxygen transport in tissue ("convection" of heat or oxygen via vessel transport and conduction of heat or diffusion of oxygen out of vessels), meaning a thermal measurement method may provide a more accurate estimation of perfusion adequacy than competing methods.

Electrolytically generated hydrogen gas sensed using a platinum electrode: has been used as an invasive local perfusion tracer in a large number of published studies, with the clearance time coefficient being related to volumetric perfusion. As the diffusion coefficient of hydrogen gas is close to that of oxygen, transport behaviour is also similar. Trans-cutaneous use of hydrogen gas as a perfusion tracer has been demonstrated in a single published study [7], but in this study the gas was introduced via an inhaled gas mixture, an impractical solution for routine monitoring. It is unclear if a self contained trans-cutaneous hydrogen clearance sensor utilising electrolytic hydrogen generation is a feasible solution.

Trans-cutaneous blood gas sensing (CO_2 and O_2): widely used clinically, indicating perfusion adequacy [8], but not volumetric perfusion directly.

Near infra-red spectroscopy (NIRS): this has also been suggested for perfusion adequacy assessment, but is currently mainly used as a research rather than a clinical tool. NIRS can only directly determine tissue oxygenation, which is not directly related to volumetric perfusion. However, perfusion adequacy is an extremely important metric clinically (e.g. see Boldt [9]). Spectroscopy in scattering media such as biological tissue is complex. Computation of the optical TPSF (Temporal Point Spread Function) is necessary for derivation of accurate absolute molarities. This involves complex and high cost optical components (usually pulsed lasers and photomultiplier tubes or avalanche photodiodes), limiting its use to research level studies. Continuous wave NIRS is a lower cost approach which can use simple LED light sources. However, it relies on heavily approximated models and has limited accuracy, although this can be improved if only differential measurements are required.

1.3 Thesis structure

This thesis is divided into a further ten chapters. Chapters two to six make up a literature review and discussion section, presenting limited experimental and simulation data. As the perfusion monitoring techniques under consideration span a wide range of disciplines, from electrochemical sensing to NMR and spectroscopy, discussion is divided between the first five chapters according to discipline.

Chapter 2 discusses thermal perfusion sensors; active, passive, and non contact. Simulation results, including of a novel active heated sensor design, and experimental studies with non contact optical heating and long wave infrared temperature monitoring presented. Controversy surrounding the relative accuracy of different bioheat transfer models is discussed.

Chapter 3 discusses application of NMR to perfusion monitoring, in particular small scale single sided permanent magnet devices. Theoretical designs for a permanent magnet NMR flowmeter is analysed, with both a simple single sided permanent magnet design and a larger device in the form of a small unit enclosing a hand or arm being studied.

Chapter 4 focusses on use of trans-cutaneous blood gas sensing in perfusion measurement, discussing application of off-the-shelf trans-cutaneous CO_2 and O_2 sensing to perfusion adequacy assessment, as well as use of H_2 as a perfusion tracer.

Chapter 5 discusses Doppler ultrasonic techniques, principally focussing on a literature review of studies of off the shelf hardware applied to capillary bed perfusion.

Chapter 6 discusses optical perfusion monitoring methods, a broad topic encompassing laser Doppler flowmetry, NIRS, pulse oximetry, capillary refill and capillaroscopy.

Chapter 7 analyses techniques from earlier chapters in term of economics and Technology Readiness Level (TRL). On these bases, NIRS, capillary refill, pulse oximetry, and laser Doppler flowmetry are selected as complimentary optical techniques for further study. Design of a modulated light system for application of the first three of these is presented. This is developed as a combination NIRS/capillary refill device in Chapter 8, and in a data logger format in Chapter 9, where development of a series of prototype pneumatic capillary refill sensors is also presented.

Chapter 10 presents results from a volunteer study comparing laser Doppler, pulse oximetry and capillary refill response to a cold challenge test. A small scale study of the capillary refill using NIRS to investigate underlying flow mechanisms is presented.

There is a final chapter (11) with conclusions and suggestions for future work.

Chapter 2

Thermal perfusion measurement

2.1 Introduction

2.1.1 Concepts

Thermal monitoring of perfusion is an idea with a long history, and many analogous devices exist in other fields of flow measurement. For example sap flow meters for use with woody plants use the transport of heat via the flowing sap as a measurement principle, and industrial flow meters designed for installation on water pipes without risk of leakage often use the heat transfer principle [10].

The reasoning behind these ideas is simple; water has a high volumetric heat capacity of approximately $4.2 \times 10^6 Jm^{-3}K^{-1}$ at 25°C. In biological tissue, the high heat capacity of blood (consisting of approximately 83% water) is complimented by a low tissue thermal conductivity, of order $0.3Wm^{-1}K^{-1}$. This low thermal conductivity further amplifies the effect of blood flow, as thermal flux from any heated "sense" element is inhibited. The effect of perfusion as a heat remover is thus more significant in the overall balance of thermal sinks.

Thermal monitoring also has apparent advantages over other competing techniques due to the nature of the thermal exchange between blood and tissue. Due to the smaller surface area to volume ratio and faster blood flow velocity, the blood flowing in larger vessels is not in thermal equilibrium with the surrounding tissue. The exact points in the cardiovascular loop at which blood reaches thermal equilibrium with its surrounding tissue is subject to debate, but the consensus would seem to be that arterial blood flowing into the periphery first reaches thermal equilibrium with the surrounding tissue at some point in the arterioles. The blood vessel diameter at which this thermal equilibrium is reached will clearly vary depending upon factors such as the tissue type and the

perfusion itself (as blood flow velocity will be a function of the perfusion), but authors have suggested that this equilibrium occurs at diameters in the range between $100\mu m$ and $1mm$ (Kutz et al [11, Chp. 2, p. 8]).

2.1.2 Comparison with competing techniques

When thermal monitoring is compared to other perfusion monitoring techniques such as laser Doppler and to a lesser extent Doppler ultrasound and nuclear magnetic resonance (NMR), it can be seen that the fact that vessels in thermal equilibria are located at the ends of the circulatory loop implies that they provide a more suitable perfusion metric than the bulk flow measurement (i.e. an average over all vessels) provided by other methods.

Doppler ultrasound allows for determination of uniaxial flow velocity, as it is a single scattering technique. Thus ultrasonic Doppler data may be processed to remove high Doppler shifts, and output a measurement of total reflectivity at low shifts that are more representative of typical capillary flow velocities. However, with only a single transmitter and receiver Doppler ultrasound will necessarily be uniaxial, so for example a large blood vessel that skirts the sensed volume in a circumferential direction may only give rise to a small Doppler shift, and thus the large cross sectional area and high flow velocity would appear as a large number of more slowly flowing capillaries. This situation would seem to reveal a fundamental problem with Doppler ultrasound, but one that could perhaps be addressed through the use of multiple transmit/receive transducers.

NMR is a technique with a large range of possible implementations, so a trade off analysis between NMR and thermal monitoring of perfusion requires consideration of specific NMR techniques. In conventional high field, high homogeneity MRI, techniques such as diffusion tensor, diffusion weighted, and diffusion spectroscopic imaging can provide detailed assessment of diffuse flow. These techniques are typically employed for assessment of interstitial fluid diffusion, allowing tissue structure to be explored in great detail non-invasively. However, blood perfusion can also be quantified using diffusion imaging. Unfortunately the field strength and homogeneity required to employ these techniques makes them wholly unsuitable for use in bedside monitoring.

Permanent magnet NMR is a development that will be discussed in a separate chapter, but the use of low strength and high inhomogeneity fields means any practical NMR based sensor for continuous monitoring will have much diminished capabilities when compared to conventional MRI. In designing a workable scheme for permanent magnet NMR, if indeed such a scheme is feasible at all, it may well be necessary to employ a processing scheme that suffers from the same large vessel susceptibility as laser Doppler

and Doppler ultrasound.

In summary it can be seen that, compared to other methods, thermal methods appear to have potential as continuous non-invasive perfusion metrics, and for that reason will be explored in more detail in this chapter.

2.1.3 Current developments in thermal perfusion monitoring

Core to periphery temperature difference is an example of one commonly used and widely accepted thermal measurement of perfusion. This has also been found to correlate with capillary refill time in numerous studies, for example Tibby et al [12], Osburn et al [13], and Lima and Bakker [14]. Kelly [15] present a review of body temperature variability studies.

However, core to periphery has many sources of error, as its sensitivity depends upon the rate of heat loss from the periphery. In real world use both the core body and ambient temperatures may be changing, and the patient may have clothes and/or bedding covering some but not all parts of their body. These factors greatly influence the skin surface temperature, reducing the usefulness of core to periphery temperature difference as a diagnostic indicator.

Hemedex Inc. manufacture an invasive thermal probe designed for measurement of perfusion [16]. A similar device is described in detail by YG Lv et al [17]. Promotional material from Hemedex and an analysis of patents granted to Henry Bowman, the founder of the company, indicate that the Hemedex probe consists of a pair of calibrated thermistors separated by approximately 10mm. In operation one thermistor is heated to maintain a constant temperature elevation above the other, and in typical use the device is implanted beneath the skin surface or within an organ, e.g. to monitor perfusion to a skin graft to determine if its blood supply is adequate.

The power required to maintain the temperature difference between the thermistors is related to perfusion, but is also a function of the tissue thermal conductivity. A model fitting algorithm is used to estimate both the tissue conductivity and the perfusion from the shape of the power versus time transient after the thermistor is first turned on. The Hemedex technique assumes that the tissue is homogeneous and that perfusion is constant during the initial warm up period (lasting approximately 1 minute). The details of the model used to convert raw data to perfusion are not disclosed, but model fitting will be discussed later in this chapter. The Hemedex device would appear to be immune to changes in tissue temperature due to its use of two thermistors. However, such a monitor would not be suitable for routine clinical use as a diagnostic indicator as

it is necessary to implant it under the skin. Hemedex market the technology for use in monitoring skin grafts, a situation where an invasive sensor is a more suitable option, as the patient has already undergone surgery.

Non invasive measurements of heat flux from both a heated or a cooled disc mounted on the skin surface have been used as perfusion metrics in numerous studies (for example a cooled sensor was used on phantom tissue by Robinson et al [18], and Li published a thesis researching non invasive heated thermistors [19]). Liu et al suggested an approach based on the thermal flux through the skin surface from thermal cycling [20]. These types of measurement are complicated by the reaction of the sympathetic nervous system to skin surface temperature. As the skin temperature rises, capillary flow tends to increase. This increase is not linear, and the rate of increase with temperature has been found to rise with increasing temperature [21]. In addition, capillary flow has been found to increase dramatically once skin surface temperature rises above core body temperature. For this reason a cooled sensor has been suggested as an approach worthy of consideration, despite the technical difficulty of such an approach (for example Cardinali constructed a prototype cooled sensor [22, p. 15]).

Dittmar and colleages have published a series of studies involving noninvasive thermal perfusion monitors attached to the skin surface [23, 24]. These involved a sensor which measured the effective thermal conductivity of the skin surface, and consisted of an inner heated disc and outer thermally conductive ring, connected via a series of thermocouples. A PID (Proportional Integral Derivative) closed loop control scheme was then used to maintain a constant 2°C temperature elevation at the central disc, and the power required to maintain this elevation recorded and converted to units of thermal conductivity. Correlation of conductivity with tissue perfusion was demonstrated, with the second device (based around a flexible polyamide PCB construction) showing particularly high sensitivity, with a small reduction in skin blood flow during a mental calculation task being clearly visible. A strong (r=0.83) correlation with laser Doppler flowmetry was demonstrated using recordings from the device attached to the hands of ten volunteers [25].

An additional complication of surface mounted sensors is heat loss from the top of the sensor. An invasive probe is fully surrounded by tissue, but when a sensor is attached to the surface, its top face must be insulated, a non trivial issue given the relatively low thermal conductivity of tissue, which is less than an order of magnitude higher than that of readily available flexible foam polymers ($0.3 W m^{-1} K^{-1}$ as opposed to $0.03 W m^{-1} K^{-1}$). In US Patent 4,859,078 Bowman et al describe a technique employing an outer cover heated to temperature equal to a disc sensor that lies in contact with the skin surface

[26]. This outer cover isolates the skin sensor from the effects of heat exchange with the surrounding environment (which primarily occurs via convection). Thermocouple array based thermal flux sensors have been proposed for use with actively cooled sensors. Here, the placement of the flux sensor between the skin surface and the actively cooled part of the sensor allows heat flux entering from the upper face of the sensor to be excluded.

The constant heat loss from the body gives rise to a heat flux through the skin. This thermal flux vector will be denoted by $\overrightarrow{\phi_q}$ for future reference. This flux may itself be a useful measure of perfusion, but it could represent a source of error in actively heated sensors, a topic that will be discussed later in this chapter.

2.2 Heat transport in biological tissue

The following review of biological heat transport models will be kept relatively brief. A more thorough overview can be found in the Standard handbook of biomedical engineering and design [11, Chapter 2].

2.2.1 Pennes bioheat equation

Pennes introduced the first widely use bioheat equation in 1948 [27]. In this model the blood is assumed to reach equilibrium with the core body temperature during passage through the heart and lungs. The arterial blood then maintains a constant temperature as it travels to the capillary bed, and upon entering the capillary bed reaches thermal equilibrium with the tissue. The capillary bed is assumed to consist of an isotropically distributed system of vessels. The blood leaves the capillary bed at a temperature equal to the local tissue, entering the veins and flowing back towards the heart in thermal isolation from the surrounding tissue. This type of model is described as a "continuum" model, as the effect of the blood vessels and blood flow upon thermal properties is treated as a bulk property of the tissue.

Despite the obvious simplicity of this model, it was found by Pennes to provide a good fit to forearm temperature profiles in a series of invasive measurements on healthy adults, and also in following studies for example Rawnsley et al [28]. Many following studies (for example Arkin et al [29]) have confirmed Pennes' hypothesis that in the larger arteries and veins blood is not in thermal equilibrium with surrounding tissue, due to the high flow velocity, high heat capacity of blood, and smaller surface area to volume ratio as blood vessel diameter increases. These assumptions remain unchanged in proposed alternative thermal models, but the effect of the blood flow in the capillary bed as a

heat transport media has been developed further.

Pennes added two terms to the thermal diffusion equation, one to represent heat exchange with the micro-vascular blood flow, and a second representing heat from metabolic processes within the tissue (q_p)

$$\rho_t C_t \frac{\partial T_t}{\partial t} = \nabla(k_t \nabla T_t) + \omega \rho_b C_b \left(T_a - T_b\right) + q_p. \tag{2.1}$$

Here C represents heat capacity, ρ is density, k_t is thermal conductivity and ω is volumetric perfusion ($ml(blood)/ml(tissue)/s$). If the blood and tissue densities and heat capacities are assumed to be approximately equal (they will both have similar values to water), then in spherical co-ordinates and assuming a uniform k_t, equation 2.1 can be expressed in the following form, which is useful for application to invasive spherical thermistors:

$$\rho C \frac{\partial T_t}{\partial t} = k_t \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial T_t}{\partial r} \right) + \omega \rho C \left(T_a - T_b\right) + q_p. \tag{2.2}$$

More recently, doubts have arisen regarding the simplifying assumptions applied to the capillary bed in the Pennes model, for example from Brinck et al [30]; the capillary blood is in thermal equilibrium with surrounding tissue for a considerable length of time as it passes through the capillaries, providing opportunity for bi-directional thermal flow both into and out of capillary vessels at different points along their path. In addition, the vessels in real tissue are not isotropically orientated, providing opportunity for structures resembling counter-current heat exchangers to arise in the matrix of capillary vessels. However, some authors have also defended the original work as it provides a good fit to the measured temperature profiles [31]. Charney [32, p. 19] provides a more in depth discussion of the debate surrounding the Pennes model, and Jiji (2009) [33] is a more up to date source covering recent developments.

2.2.2 Wulff continuum model

Wulff in 1974 [34] was one of the first authors to explore alternatives to the Pennes model. By treating the blood flow as an enthalpy transport mechanism, the thermal flux q at a point in the tissue can be seen to be equal to:

$$q = k_t \nabla T_t + \rho_b h_b v_b, \tag{2.3}$$

where h_b is the specific enthalpy of the blood, and v_b the mean blood flow velocity (volumetrically weighted). Assuming that the micro-vascular circulation is in constant

thermal equilibrium with the surrounding tissue , and the divergence of $\rho_b v_b$ is zero (i.e. there are no sources or sinks of blood within the tissue), Wulff then derived the following bioheat equation:

$$\rho_t C_t \frac{\partial T_t}{\partial t} = k_t \nabla^2 T_t + \rho_b C_b \mathbf{v}_b . \nabla T_t + q_p. \tag{2.4}$$

Note that in deriving equation 2.4, k_t is assumed to be constant throughout the tissue, and a q_p term has been introduced to account for metabolic heat generation. Comparing equations 2.1 and 2.4, it can be seen that the term representing heat flux from inflowing arterial blood reaching thermal equilibrium with the surrounding tissue has been replaced by a term representing heat transport via flowing blood. Where Pennes assumed the blood flow acts to transport heat to the capillary bed, and arterial and venous flows are not thermally significant to the tissue they flow through, Wulff assumes that all flowing blood is in continual thermal equilibrium, and acting as a heat transport agent at a local level. Hence in the Wulff model there is no large body scale heat transport mechanism, only local scale transport.

Klinger [35] introduced a very similar model in 1974, only with a more general definition of the mean blood flow velocity, allowing for the velocity to be both spatially and temporally variant.

2.2.3 Chen-Holmes continuum model

In 1980 Chen and Holmes [36] developed a more complex model, incorporating the ideas of Wulff and Klinger together with the Pennes equation term for heat flux from equilibrating arterial blood. Chen and Holmes gave the thermal flux as:

$$q = k_t \nabla T_t + \rho_b C_b \omega_b^* (T_{art}^* - T_t) - \rho_b C_b \mathbf{v} . \nabla T_t + \nabla . k_p \nabla T_t. \tag{2.5}$$

It can be seen that this equation includes a term similar to Pennes, but the perfusion ω^* and arterial temperature T_{art}^* are now variables, taking different values dependent upon the specific volume being modelled, this being denoted by the superscript.

The third term in equation 2.5 can be seen to be similar to the heat transport term from Wulff and Klinger, with the vector $\mathbf{v}_b$ being equal to the volumetric blood flow vector per unit area. The fourth term was introduced by Chen and Holmes to account for blood flow that is in near equilibrium with surrounding tissue. The new variable k_p being a conductivity like effect from the flowing blood, a complex effect related to the thermal equilibrium length of the blood vessels, the distribution of flow vectors relative

to the thermal gradient, and the area number density of vessels.

In practice the complex nature of the k_p limits the usefulness of Chen-Homes, since detailed knowledge of the vascular network, such as the distribution of vessel diameters and lengths is required. The Chen-Holmes model has also been criticised for mis-modelling the effect of closely spaced arteriole vein pairs (e.g. by Kreith [37]).

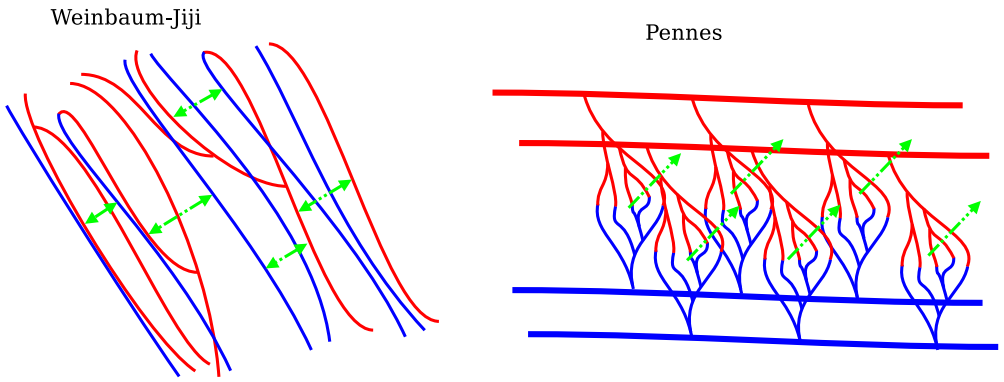


Figure 2.1: Comparison of the Weinbaum-Jiji heat transfer hypothesis to Pennes model, arrows show heat flux.

2.2.4 Weinbaum, Jiji and Lemons bioheat model

In 1984, Weinbaum, Jiji and Lemons published [38],[39] a model including counter-current heat exchange between arteriole-vein pairs. This was a development of a model proposed by the same authors in 1979 [40], which consisted of artery-vein pairs as two parallel cylinders of equal diameters with balanced bleed-off in the plane normal to the cylinders. A model which was based extensively on studies of rabbit limbs performed by Weinbaum et al.

In their 1984 model, Weinbaum et al broke the outer 1cm of tissue down into three layers; a deep, intermediate, and cutaneous layer. In the deep tissue they proposed a system of three coupled equations:

$$\rho_b C_b \pi r^2 \overline{\mathbf{v}}.\frac{dT_{art}}{ds} = -q_{art}, \tag{2.6}$$

$$\rho_b C_b \pi r^2 \overline{\mathbf{v}}.\frac{dT_v}{ds} = -q_v, \tag{2.7}$$

$$\rho_t C_t \frac{\partial T_t}{\partial t} = \nabla.k_t \nabla T_t + \left(ng\rho_p C_p (T_{art} - T_v) - \rho_b C_b n \pi r^2 \overline{\mathbf{v}}.\frac{d(T_{art} - T_v)}{ds} + q_p \right). \tag{2.8}$$

Where q_{art} is the heat loss from the arteries, and q_v is heat gain by the veins. T_{art} and T_v are the mean arterial and venous blood temperatures, and $\overline{\mathbf{v}}$ is the mean blood flow velocity in the vessels. There is a perfusion bleed off term g, and a weighted blood flow velocity per vessel surface area s. The first two equations 2.6 and 2.7 describe heat conduction into and out of the tissue from veins and arteries. The third equation 2.8 describes the perfused tissue. It can be seen that the second term in this equation is similar the Pennes equation perfusion term, only with a mass term g and vessel number density term n used in place of a volumetric perfusion.

The model assumes that blood exits the capillary bed at the venous temperature T_v, an assumption that has been criticised (for example by Kreith [37]). The model is also difficult to implement, and does not have unique solutions. It is not possible to uniquely determine T_{art}, T_v, and T_i without simplifying assumptions. For this reason in 1985 Weinbaum and Jiji derived a simplified model.

2.2.5 Simplified Weinbaum-Jiji model

In their 1985 model [41], Weinbaum and Jiji made two simplifying assumptions:

1. The tissue temperature T_t is the average of arterial and venous temperatures:

$$T_t \approx \frac{T_{art} + T_v}{2}. \tag{2.9}$$

2. Heat conducted from an artery is mostly conducted to a paired vein:

$$q_v \approx q_{art} \approx \sigma_{\Delta} k_t (T_{art} - T_v). \tag{2.10}$$

σ_{Δ} is a geometry factor representing the thermal resistance between two vessels embedded in tissue. Chato derived the following for σ_{Δ}, based on an assumption of parallel cylindrical vessels (radius r, separation l) with uniform surface temperatures embedded in an infinite tissue slab

$$\sigma_{\Delta} = \frac{\pi}{cosh^{-1}(l/2r)}. \tag{2.11}$$

By substituting their two assumptions 2.9 and 2.10 into the previous Weinbaum-Jiji bioheat equation, the following equation can be derived:

$$\rho_t C_t \frac{\partial T_t}{\partial t} = \nabla . k_{eff} \nabla T_t + q_p. \tag{2.12}$$

Here a new "effective" thermal conductivity has been introduced, k_{eff}, taking into account both thermal conduction through the tissue and conduction like effects from blood flow:

$$k_{eff} = k_t(1 + Pe_i V(\xi)). \tag{2.13}$$

ξ is a dimensionless distance, $\xi = x/L$, the ratio of the distance into the tissue layer (x) relative to the tissue layer thickness (L). $V(\xi)$ is a function of the vascular geometry. Pe_i is the thermal Péclet number for the inlet vessel to the tissue slab, defined as:

$$Pe_i = \frac{2\rho_b C_b r_i v_i}{k_b}. \tag{2.14}$$

Where subscript i denotes the inlet vessel into the tissue slab extending from 0 to x.

Despite the two simplifications made, the simplified Weinbaum-Jiji still requires knowledge of the vascular structure for the function $V(\xi)$ to be derived. This stands in contrast to the Pennes equation where only a single parameter of perfusion is required.

2.3 Numerical modelling of the bioheat equation

In order to further assess the feasibility of a thermal perfusion sensor a numerical model of perfused tissue would be extremely useful. Unfortunately the question of which equation to use is not a simple one, as illustrated by the preceding review of bioheat equations.

2.3.1 Choosing a model

Despite the fact that its validity has been widely questioned, the Pennes model was used in the following numerical simulations. The decision to use this model was made for a number of reasons:

1. Pennes has been widely used by many authors, and is perhaps the most widely used bioheat equation.

2. The Pennes model is easy to apply to different tissue types and perfusions, taking only tissue conductivity and volumetric perfusion (in ml/ml/s units) as inputs.

3. The main alternatives to Pennes (Chen-Holmes and Weinbaum-Jiji) both contain similar terms to the Pennes equation. Specifically the second term on the right of

equations 2.5 and 2.8. This suggests that in real tissue, the heat exchange will be similar to that modelled by the Pennes equation if the tissue thermal conductivity in the Pennes model is made to increase with blood flow.

As equation 2.13 illustrates, any effective thermal conductivity term will be a complex factor of blood vessel geometry. For this reason no attempt has been made to move beyond the basic Pennes equation. However, it should be noted that if several simulations are made using the Pennes equation with differing tissue thermal conductivities, then these can be used to gauge how the mis-modelling of inter-vessel heat flux in Pennes effects the results. The real world behaviour may strongly resemble an interpolation between simulations with different tissue thermal conductivities, with the interpolated conductivity increasing with perfusion. Wren et al describe such an approach in [42].

Having determined which model to use, the number of dimensions to use in the model needs to be determined. Clearly real world tissue is three dimensional, but by making symmetry arguments this can be reduced, so a spherically symmetric sensor (e.g. an invasive thermistor as in the Hemedex sensor) would only require a single radial dimension. If a circular heated disc of uniform temperature is placed on the skin surface, then a two dimensional model can be used in cylindrical co-ordinates. The assumption of a sensor of negligible heat capacity and uniform temperature can be seen to be a valid approximation by considering the properties of typical materials that may be used. For example copper foil; copper has a heat capacity of $3.5J(ml)^{-1}K^{-1}$, slightly lower than that of typical tissue, and a thermal conductivity of $400Wm^{-1}K^{-1}$, more than three orders of magnitude higher than typical tissue, thus a thin foil could be used.

The Pennes equation in spherical co-ordinates is shown in 2.2. In cylindrical co-ordinates, and assuming cylindrical symmetry and homogeneous k_t, the equation is:

$$\rho C \frac{\partial T_t}{\partial t} = k_t \frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial T_t}{\partial r} \right) + k_t \frac{\partial^2 T_t}{\partial z^2} + \omega \rho C \left(T_a - T_b \right) + q_p. \tag{2.15}$$

A one dimensional model is clearly easier to evaluate numerically, but may not be applicable to a non invasive disc. In order to assess the size of the error introduced by treating a surface disc as a hemisphere, two numerical simulations were conducted with identical tissue thermal conductivity of $0.2Wm^{-1}K^{-1}$ and changing perfusion. The simulated sensor was maintained at a constant 1 degree temperature elevation above the tissue, and the input power recorded once thermal equilibrium was reached between the sensor and tissue.

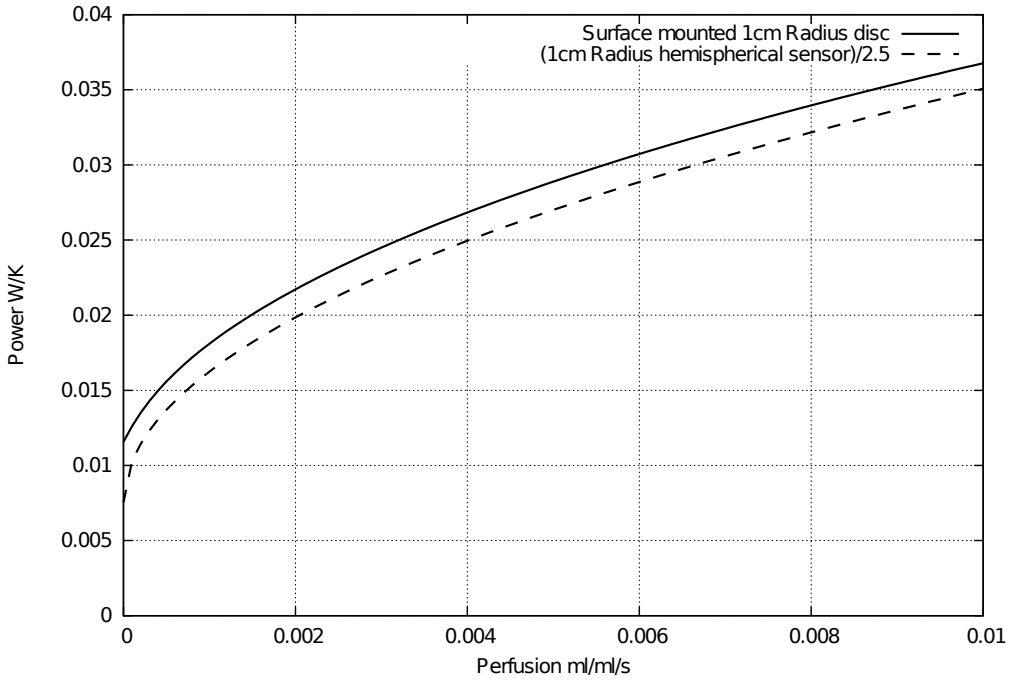


Figure 2.2: Comparison of a hemispherical sensor simulated in spherical co-ordinates to a disc sensor in cylindrical co-ordinates after running to equlibrium

From figure 2.2 it can be seen that a hemispherical sensor has approximately 2.5 times the thermal flux of a disc sensor at the same perfusion (the hemispherical dataset is scaled in the plot). This might be expected given the larger surface area of $2\pi r^2$ as opposed to πr^2, and the fact that a hemisphere is a more effective shape for heat loss via conduction than a disc. For analysis of sensor performance a linear scaling factor is easily accounted for when interpreting model results, but a closer inspection of the graph at low perfusion levels reveals a non-linear scaling between the two models at low perfusion.

For this reason it was determined that a two dimensional simulation of the Pennes model in cylindrical co-ordinates with cylindrical symmetry about the ϕ axis would be required to give a reasonable assessment of non-invasive sensor performance.

2.3.2 Simulation of a heated disc sensor

The numerical simulation of equation 2.15 (written in the GNU-Octave programming language) was used to assess the practicality of a 1cm diameter disc sensor. The simulation was run for 100 simulated minutes over a range of different tissue perfusions and conductivities. By inspection of the data it was found that 100 minutes was enough time for the disc to reach pseudo equilibrium with the tissue in all operating conditions (as the tissue will never reach full equilibrium, "pseudo" equilibrium was defined as a

less than 4% change in input power over the final 10 simulated minutes). Typical values for the thermal properties are based on those of Bowman et al [43].

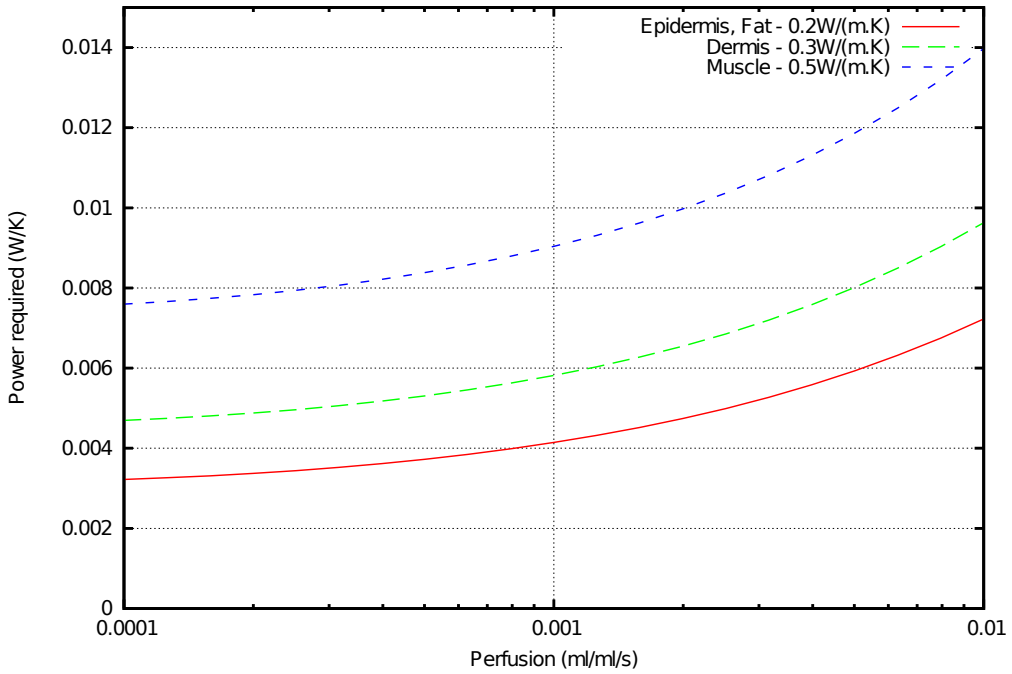


Figure 2.3: Comparison of a simulated 1cm diameter disc sensor run to equilibrium under a range of operating conditions.

The key in figure 2.3 gives examples of tissue types with typical thermal conductivities equal to the simulated values. It can be seen there is a clear increase in power input to the sensor with increasing perfusion for all the tissue types. As before the sensor was held at a constant temperature elevation relative to the tissue. Typical perfusion values for healthy muscle and skin tissue are of order 0.001ml/ml/s, so changes both above and below the typical baseline can be detected as input power fluctuations.

However, if absolute perfusion is required, rather than merely perfusion trends over time, then the sensitivity to tissue thermal conductivity is clearly a serious problem. It can be seen that at the highest thermal conductivity and lowest perfusion, input power is still higher than lowest thermal conductivity and highest perfusion (0.01ml/ml/s is around the highest value for volumetric perfusion typically seen in human skin, muscle or subcutaneous fat tissue).

A further problem that would be faced when attempting to use a simple heated disc is uncertainties in the unperturbed skin surface temperature. In any practical sensor design, a second temperature sensor physically separated from the heated disc would be required in order to measure the skin surface temperature as it changes under the effect of environmental conditions and fluctuating perfusion and core body temperature. Designing a sensor that was isolated from the surrounding environment in as similar

a fashion to the heated disc as possible would be difficult, and whilst increasing the temperature difference between the heated disc and surface temperature sensor would reduce the sensitivity to environmental conditions, it would increase the potential for bio-regulatory processes to increase perfusion under the heated area and thus bias the measurements to an unpredictable extent. Li developed and tested such a sensor, with some limited success [19, p. 71]. Li's device apparently suffered severely from skin contact issues, a problem that could be overcome with a more flexible thin film self adhesive sensor (although Li did not suggest this solution in the publication).

Poor skin contact is a further source of error with non-invasive skin surface thermal sensors. To investigate this, the effect of a contact gap consisting of a 0.1mm layer of static, non convecting air between the sensor surface and skin was simulated using Pennes's bioheat equation, and the total thermal flux from the sensor under constant temperature elevation was measured with and without the air gap present. The percentage error in input power due to the air gap was calculated and is plotted in figure 2.4

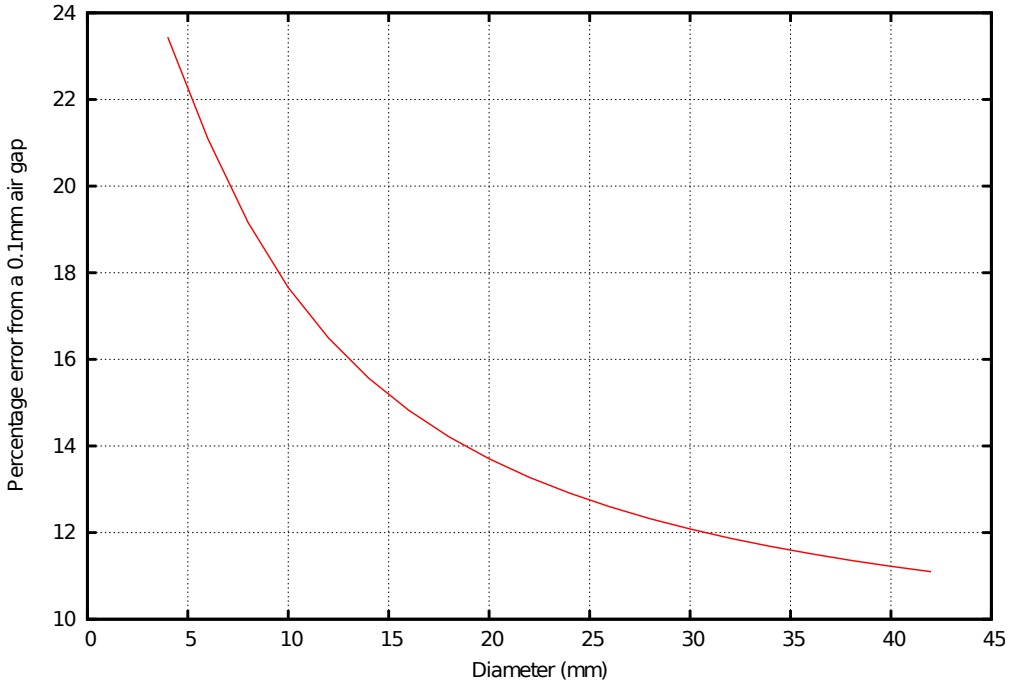


Figure 2.4: Effect of contact error on thermal disc sensors of differing diameters run to equilibrium. Perfusion was 0.001ml/ml/s, conductivity $0.3 W m^{-1} K^{-1}$.

It can be seen that a larger sensor results in a less significant contact error. However, it would appear from the graph that as the sensor diameter tends to infinity the percentage error does not tend to zero. Running a single simulation of an extremely large 50cm diameter sensor confirms this, the error in input power was found to stabilise at around 10% at large sensor sizes. The reason for these effects is obvious when the Pennes

equation is considered; at small sensor sizes, the conduction loss becomes larger per unit surface area of sensor, so the thermal impedance of the air gap becomes more significant. At very large sensor sizes, the conduction losses become insignificant compared to the perfusion loss, and the temperature distribution becomes a one dimensional function of depth only, with a uniform vertical temperature gradient below the sensor surface. The heat flux per unit sensor area tends to a function of the product of perfusion and tissue thermal conductivity, with the edge effects around the sensor becoming increasingly insignificant, and thus the air gap induced error tends to a constant value as the diameter tends to infinity.

Referring back to figure 2.3, it can be seen that the 18% gap induced error on a 1cm diameter disc is significant enough to cause a perfusion of 0.001 to appear as 0.0001ml/ml/s. This is clearly significant, but a 0.1mm air gap under the sensor is perhaps quite a pessimistic model of contact error, particularly for use on neonates where the skin is thin and there is little body hair. If we consider a realistic sensor consisting of a poly-amide flexible PCB with copper cladding and a thin bio-compatible coating and pressure sensitive adhesive layer, then as the sensor conforms to the skin surface the thin layer of pressure sensitive adhesive will be able to maintain a consistent bond to the skin surface, and so air gaps should be insignificant providing the skin surface is smooth and the sensor is applied with care. The skin contact problem could be compared to the contact problem in trans-cutaneous oxygen monitoring, where even a small gap can result in very poor measurement accuracy, yet trans-cutaneous oxygen sensors are widely and successfully used in clinical practice.

Optical heating and mid infra-red non contact temperature monitoring is a technique that eliminates contact error entirely. An experimental fibre optic heating system is described later in this chapter.

In order to better understand the heat flow out of the disc sensor, the steady state temperature distribution below a 1cm radius disc maintained at a uniform 1°C elevation above baseline was plotted for a typical perfused and unperfused state. The tissue thermal conductivity was unchanged between the two simulations at $0.2 W m^{-1} K^{-1}$.

In figures 2.5 and 2.6 temperature contours at 0.1°C degree increments are plotted in the cross sectional plane. Depth into the tissue in mm units from the surface is in the vertical axis. It is immediately apparent that there is a significant difference between the two plots; the presence of perfusing blood acts to limit the spread of heat into the tissue, and in the perfused case the final contour at 0.1°C elevation above baseline temperature is much closer to the heated sensor. A more subtle change is also evident in the geometry of the heat distribution; the perfused state causes the temperature contours to bunch

more tightly under the heated disc.

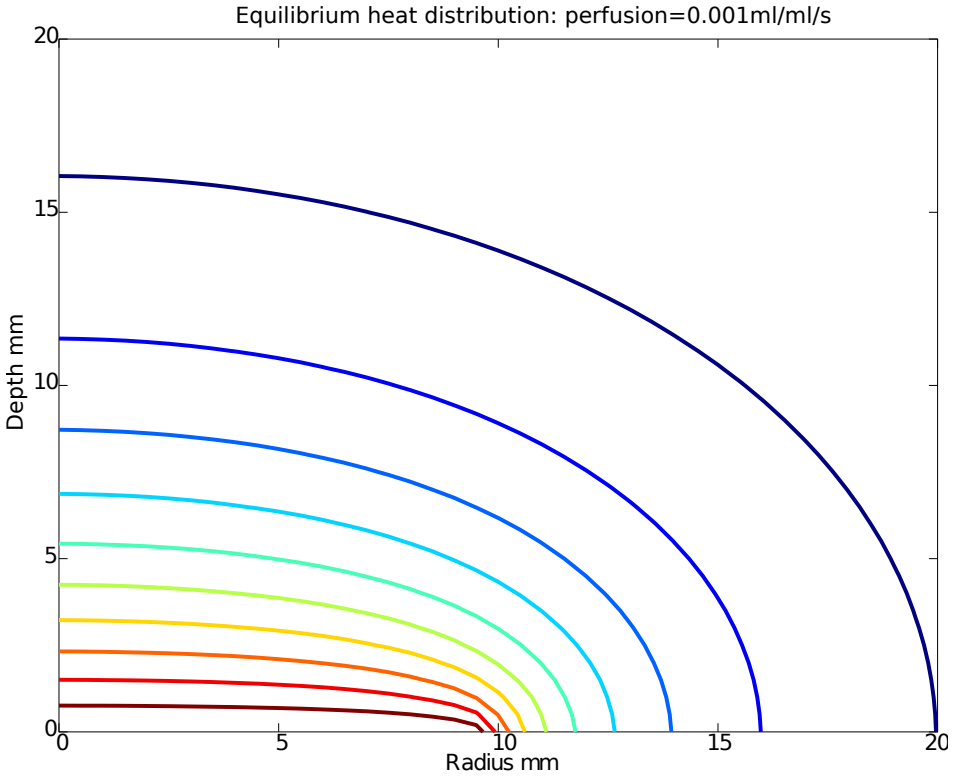


Figure 2.5: Temperature contours beneath a thermal disc sensor in a perfused state.

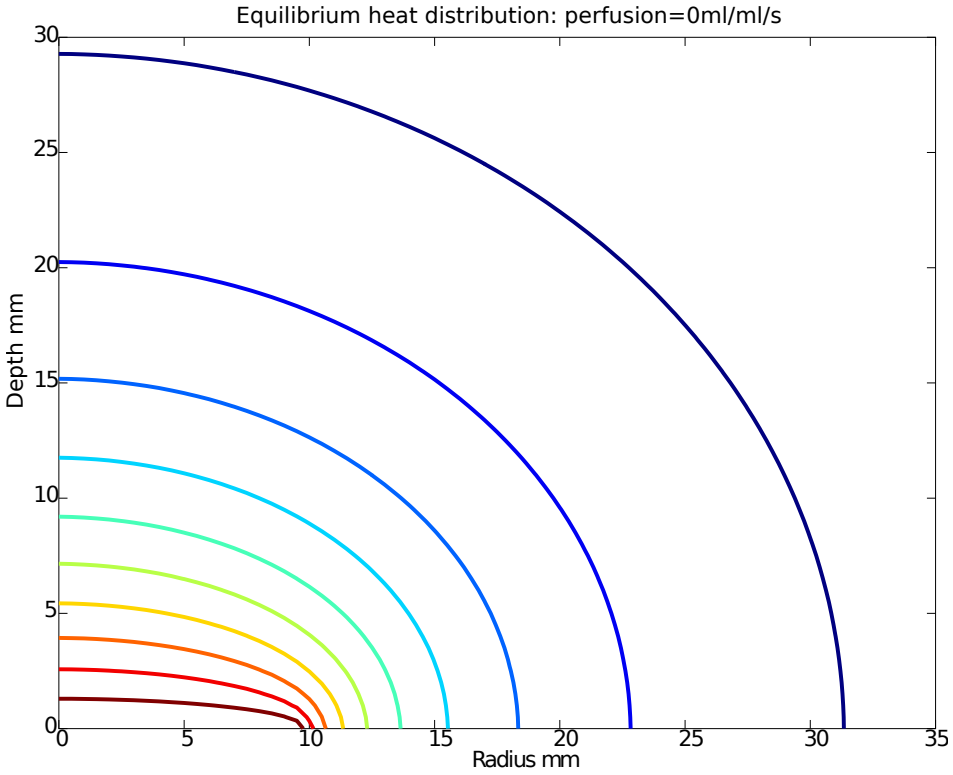


Figure 2.6: Temperature contours beneath a thermal disc sensor in an unperfused state.

2.3.3 Designing a more useful sensor

The differing heat distribution hints at an interesting property of the Pennes model: the perfusion heat flux term is fundamentally different in behaviour to thermal conduction. From Fourier's law of heat conduction it follows that any system involving only conduction of heat between fixed temperature sources and sinks of thermal flux has a temperature distribution that is invariant under changes in thermal conductivity, and possesses an internal thermal flux that scales linearly with thermal conductivity. By implication it follows that a similar effect should be seen for both the Weinbaum-Jiji and Chen-Holmes bioheat equations, due to similar "thermal washout" terms in the equations.

Thus if a sensor can be designed such that a change in the temperature distribution caused by the Pennes perfusion term can be detected, then it will be resistant to offset effects caused by variable or unknown thermal conductivity. Unfortunately due to the high volumetric heat capacity of tissue, any measurement approach that directly relies upon changes in tissue temperature will have a slow response time. One simple way to reduce response time would be to make use of the fact that heat conduction is a linear phenomenon (again from Fourier's law). Thus if a disc and outer concentric ring are heated to the same temperature elevation above baseline, then the ratio of the input powers to the disc and ring required to maintain equal and constant temperature elevations will be a constant in the absence of any perfusion. An analysis of the earlier simulation contour plots will show that in the presence of perfusion, heat flux in the central disc will increase relative to that in the outer ring (as the contours bunch more tightly).

Such a sensor would allow a zero perfusion state to be detected unequivocally without knowledge of the tissue thermal conductivity. The response time would also be relatively fast, as the output would be based upon input power ratio rather than any function of absolute temperature. When the disc simulation was repeated for different tissue thermal conductivities, and the ratio of input powers calculated for an inner disc and outer ring, it was found that this ratio was a function of the product of perfusion and conductivity, meaning that an absolute perfusion value could not be obtained without knowledge of the tissue thermal conductivity. However, tissue thermal conductivity can be estimated from the total input power, so use of a model fitting algorithm to estimate the tissue conductivity would allow absolute perfusion to be estimated. With a concentric ring sensor any model fitting errors would result in perfusion sensitivity errors rather than the perfusion offset error that would be seen with a simple heated disc sensor. Cardinali demonstrates a similar model fitting algorithm in [22, p. 20].

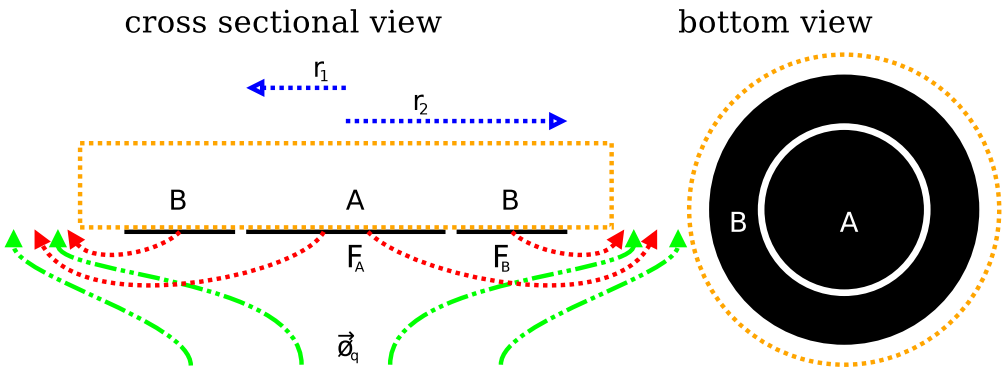


Figure 2.7: A proposed thermal perfusion sensor consisting of a heated disc and outer concentric ring.

Figure 2.7 shows a diagram of the proposed concentric ring sensor design. The inner disc is labelled 'A' and the outer ring is 'B'. An outer cover of insulator is shown in dotted orange. This serves to isolate the sensor from highly variable environmental conditions such as air currents etc.

One source of error that will still be significant in the concentric sensor is the effect of upwelling thermal flux from within the body $\overrightarrow{\phi_q}$. As this flux meets the insulated sensor it will cause a buildup of heat in the sensor centre, resulting in a decrease in input power to the central disc, and a bias in the power ratio. It may be possible to correct for some of this effect, as the upwelling thermal flux will be proportional to perfusion, and the heat buildup in the centre of the sensor will be related to the flux and thermal conductivity of the tissue.

Figure 2.8 shows a simulation of a 1cm radius sensor consisting of a 5mm radius inner disc and a 10mm outer radius concentric ring. It can be seen that the response time of the ratio metric is considerably lower, both at turn on and especially following a step change in perfusion. The time after which a change in perfusion can be distinguished from the tissue warm up transient is also much shorter with the concentric sensor. The total power sensor requires 5 minutes to warm up the surrounding tissue before the input power in the unperfused state drops below the steady state value for 0.001ml/ml/s perfusion, whereas the equivalent time for the concentric sensor is under 3 minutes.

It should be noted that Bayesian techniques such as Kalman filtering can be used to reduce the response time of thermal sensors. Use of such techniques has not been explored in this work, but similar techniques have been used to process thermal sensor data by other authors (for example by Potocki [44]). Any filtering technique that reduces the sensor response time would tend to amplify any small fluctuations, such as those caused by changes in core body temperature or changes in skin surface temperature induced by changing environmental conditions. Thus the extent to which the response time can

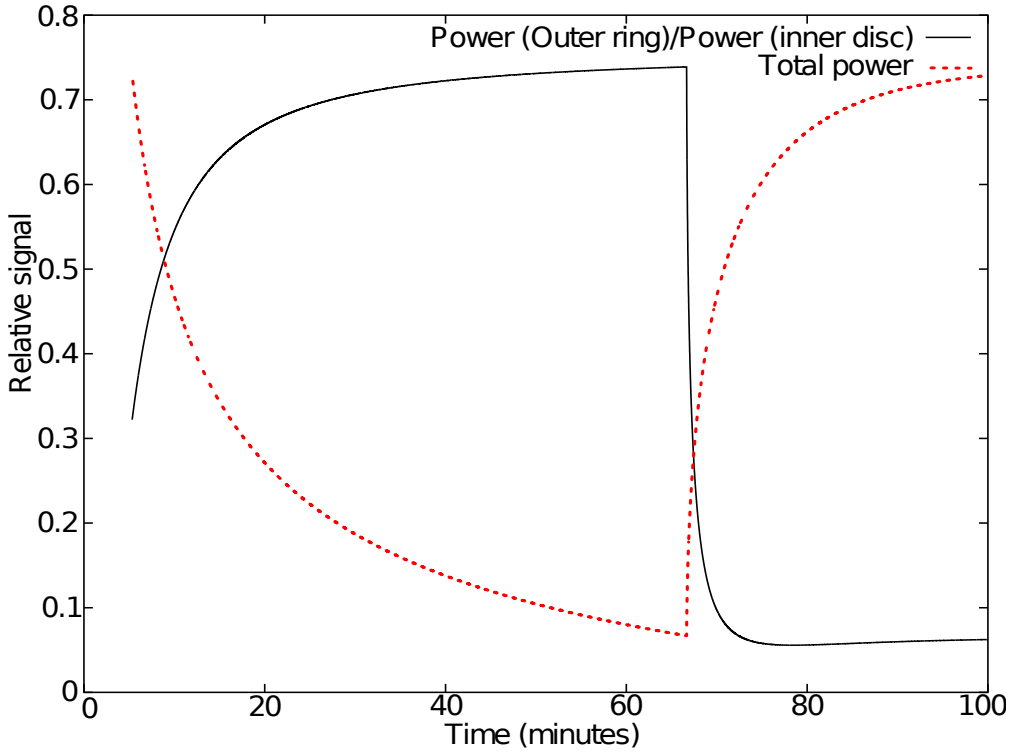


Figure 2.8: Simulation of the proposed heated disc and outer concentric ring sensor. The perfusion was stepped from 0 to 0.001ml/ml/s at 67 minutes. Tissue thermal conductivity was $0.2 W m^{-1} K^{-1}$.

be reduced via such methods will be limited in practice. Use of a non-linear Extended Kalman filter, or perhaps more advanced techniques such as Unscented Kalman filtering or particle filtering to fit a model to the sensor turn on transient may be an effective way to estimate both initial perfusion and thermal conductivity of the tissue, but this area has not been explored by the author. The extent to which fitting to the warm up transient would help reduce uncertainty in the tissue thermal conductivity is limited by the fact that the perfusion will be perturbed by the sensor itself, due to the sympathetic nervous system response to the sensor heating. This will result in an unknown perfusion perturbation as a function of time as the body reacts to the temperature perturbation over a similar time-scale to the warm up transient (5 to 10 minutes), and will make any model fitting inaccurate, even if the model is itself an accurate representation.

Figure 2.9 shows a simulation of the same 1cm outer radius sensor with the heating power ratio plotted at varying perfusions. It can be seen that in the unperfused state, the output is invariant under changes in tissue thermal conductivity. The sensor also has a logarithmic response to tissue perfusion, a useful characteristic for measurement of biophysical variables.

Figure 2.10 illustrates a possible application of the heating technique described in US

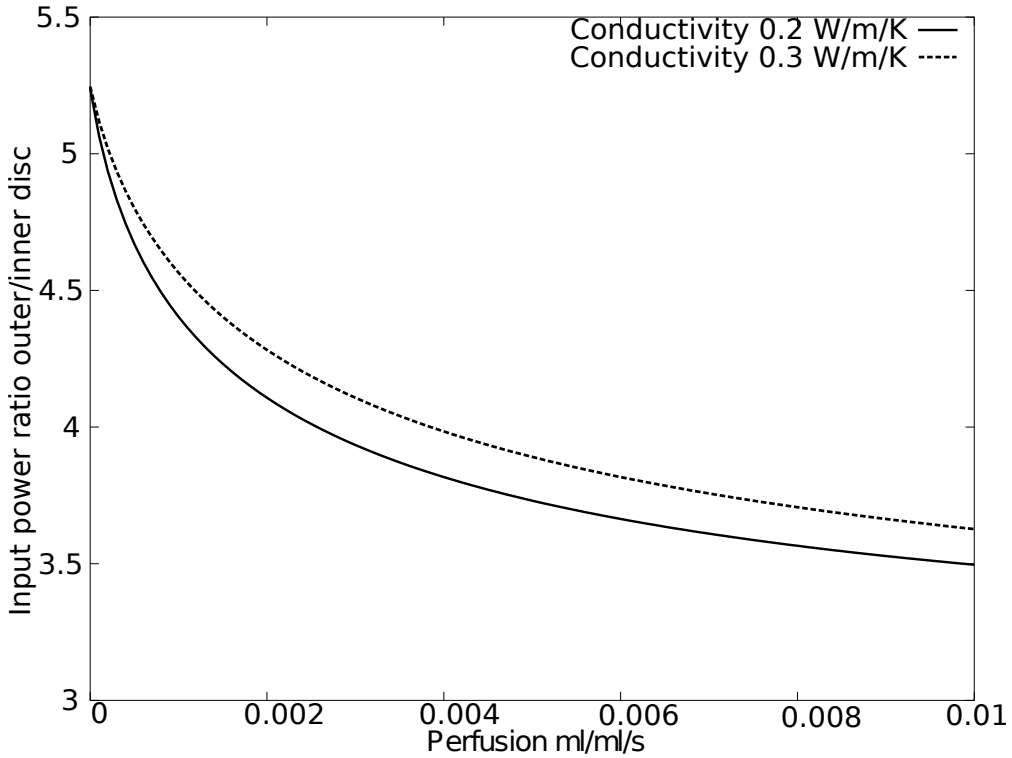


Figure 2.9: Simulation of the proposed heated disc and outer concentric ring sensor run to equilibrium at varying perfusion

patent 4,859,078 [26] to a concentric ring sensor. A covering of copper clad poly-amide flexible PCB over the sensor serves to isolate the sensor from changing environmental conditions ('C' in the diagram). A separate thermistor and feedback loop could be used to maintain the top of the sensor at the same temperature as the lower surface. This would reduce any flux through the sensor (F_L) to zero.

Just like a more basic heated disc sensor, a concentric ring sensor would require heating to constant temperature elevation above baseline. As previously discussed, achieving a consistent temperature elevation is problematic, and would require a second, remote thermistor to be used with a basic disc sensor. However, the total input power to the concentric ring sensor will be equivalent to the power input of a heated single disc with the same outer radius. Thus it can be seen that total power could be used as a second measurement that varies linearly with temperature elevation above baseline skin surface temperature (when in equilibrium with the tissue), as well as being a more complex function of perfusion and tissue thermal conductivity (as previously investigated).

Unfortunately a simple approach using an outer loop to control the temperature setpoint and maintain constant total input power would not give good results under changing perfusion. As the perfusion increases, the input power for a given temperature elevation will also increase, and so if a constant total power is maintained, the temperature

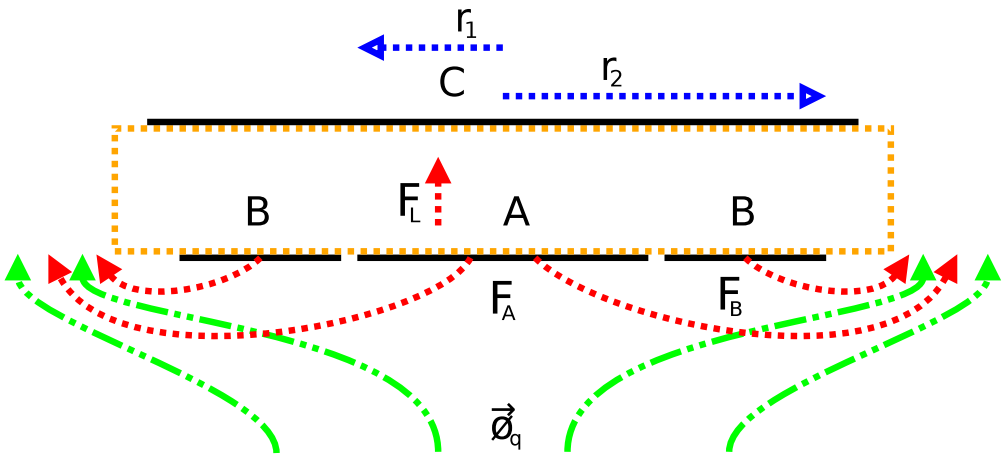


Figure 2.10: A top sensor layer heated to a temperature equal to the lower skin contact layers would dramatically reduce the effect of environmental disturbances.

elevation will rise and fall with changing perfusion, biasing any ratio output.

In order to maintain an approximately constant temperature elevation above baseline, it may be necessary to use some form of continuous model fitting in order to distinguish changing perfusion from changing skin baseline temperature. For example an Extended Kalman filter would be well suited to online model fitting and optimal state estimation.

An alternative simpler approach might be to use the input power ratio together with a calibration curve (similar to that in figure 2.9) to convert measured ratios to volumetric perfusion. The calibration curve could be derived using model fitting to the warm up transient. Once the perfusion has been derived from the power ratio, a second calibration curve, generated from the model fit, and similar to figure 2.3 could be used to estimate the temperature offset above baseline from the total power input.

The temperature elevation could then be regulated to maintain a consistent offset of perhaps 3°C, a higher elevation giving smaller artefacts from changing conditions, but with more risk of bias from increased perfusion induced by the heating.

2.4 Passive thermal flux sensing

The thermal flux through the skin surface $\overrightarrow{\phi_q}$ has previously been discussed as a source of error. However, if the assumption is made of negligible metabolic heat flux q_m, then $\overrightarrow{\phi_q}$ can be seen to be related to perfusion and the temperature elevation of the core body above ambient, and so there is potential for this flux to be used as a perfusion measurement, particularly if a patient's core body temperature can be continuously monitored (e.g. via an infra-red non contact ear probe).

For example, if a thermal flux sensor was placed on a patients limb and the limb diameter at the attachment site was recorded, then a dataset consisting of thermal flux, skin surface temperature and core body temperature versus time would allow a 2D thermal simulation of the limb cross-section to be fitted to the data. This would however require assumptions about the distribution of perfusion within the limb, and the tissue thermal conductivity would need to be known. It may prove feasible to infer the tissue conductivity from the effect of changing skin surface temperature on flux, but this has not been investigated in simulation.

Thin film thermal flux sensors are widely available commercially. These are based on flexible poly-amide PCB material with a large number of opposed thermocouple junctions on top and bottom sides. When a thermal flux passes through the film, the thermal resistance creates a small temperature differential, typically of order milliKelvin. The opposed junctions convert this to a millivolt level output proportional to the flux (typically there are several hundred junctions in series). Many also contain a separate junction that can be used to monitor device temperature.

It would be tempting to simply coat one side of a commercial sensor with adhesive, so it could be stuck to the patients skin. Unfortunately this is unlikely to work well, as the upper surface of the flux sensor may have dissimilar properties to the patients skin. The author proposes a novel passive sensor design as one possible solution to this problem.

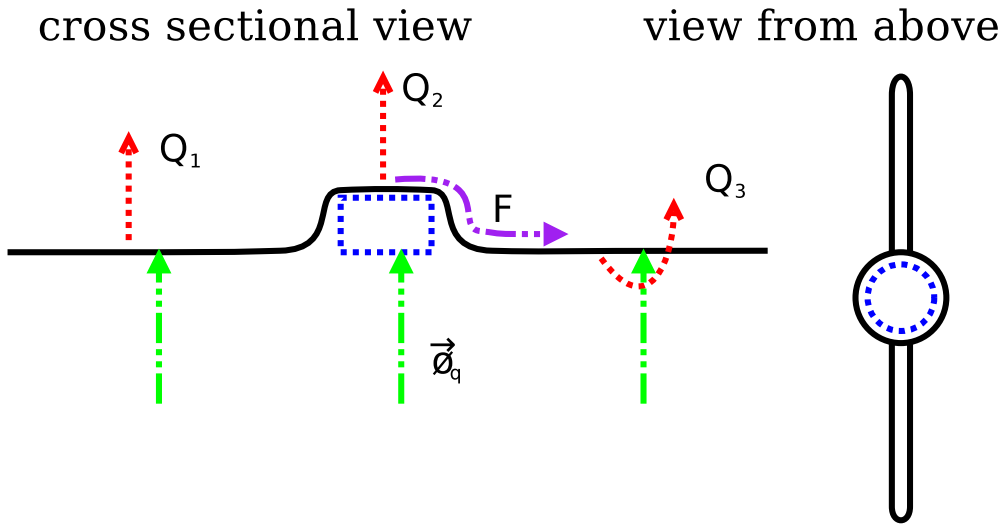


Figure 2.11: A novel passive sensor may allow skin thermal flux to be measured.

Figure 2.11 shows a diagrammatic view of the proposed sensor. The rate of heat loss from the top of the adhesive sensor is shown as Q_1 and Q_2. This would not necessarily be equal to the skin thermal flux intercepted by the sensor. However, the much higher edge length to surface area ratio of the two "wings" that spread out from the central disc region of the sensor would make the thermal resistance from the "wings" to the

surrounding skin surface much lower than that from the sensors circular central area to surrounding skin surface. Thus any mismatch in heat loss from the sensor top, shown as F on the diagram, would flow through the wings and pass into the surrounding skin, shown as Q_3.

Hence it can be seen that the thermal flux through the flux sensor (shown as blue dotted outline) is roughly equal to the equilibrium skin flux $\overrightarrow{\phi_q}$, even if the sensor top surface is not losing heat to quite the same extent as the skin surface (e.g. the skin may be sweaty, leading to very significant evaporative losses). The "wings" would also serve as attachment points for flat flexible cable connections to the bedside monitoring equipment, as any thermal flux from the cabling would be isolated by the low thermal impedance. The vertical scale is exaggerated in the diagram, and the flux sensor could use the same thin film construction as commercial thermal flux sensors, meaning the entire sensor could be a disposable adhesive thin film device with a total thickness in the range of a few hundred microns.

2.5 Experiments in non contact thermal sensing

Although it is debatable to what extent variable skin contact is an issue effecting thermal sensor error, non contact optical heating and temperature monitoring was attempted in the laboratory to assess the usability of recently available cheap (<$5) integrated mid and long wave infra-red thermopile sensors. For example the Texas Instruments TMP006 [45]. Melexis semiconductor also market a range of similar sensors, for example the MLX90616 [46], but these are larger metal cased designs, and so unsuitable for use in thin film flexible sensors compared to the BGA (Ball Grid Array) mounted and flip chip packaged TMP006. A TMP006 evaluation board was used in this experiment.

Figure 2.12 shows a schematic of the experiment. To read device temperature and target temperature, the TMP006 sensor was polled at 4Hz over an I^2C bus from a Texas Instruments development board connected to a PC using USB. A high power fibre coupled 630nm LED was used for non contact optical heating, this being driven with a constant current LED driver, and the target was black ABS plastic, as it was deemed too risky to attempt

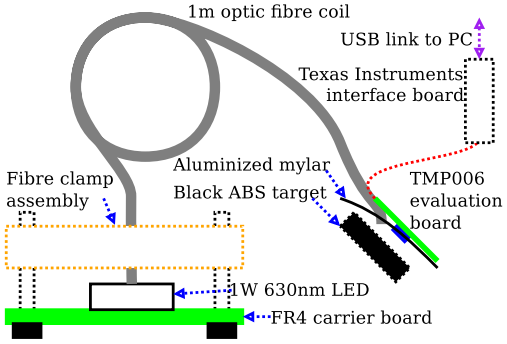


Figure 2.12: Diagrammatical view of the experimental non contact heating apparatus.

measurements on a volunteer with a highly collimated high intensity light source.

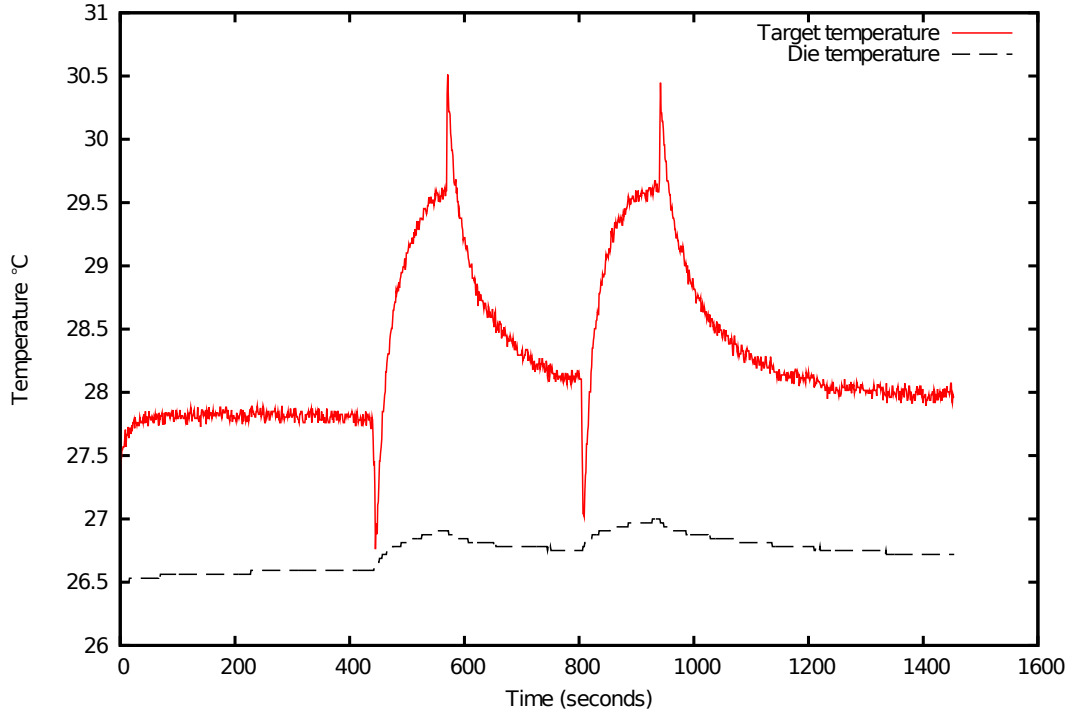


Figure 2.13: Raw results from the non contact heater.

Raw data from the sensor is graphed in figure 2.13. The optical heating is turned on from 420 to 580 seconds and 800 to 950 seconds. It can be seen that there are large spikes in the raw target temperature data when the optical illumination is turned on or off. This behaviour matches that described in the datasheet - rapid device temperature changes result in temperature gradients across the sensor die that bias the target temperature. The device datasheet [45] describes a simple finite impulse response filter which may be used to correct the recorded target temperature using the past 4 device temperature readings obtained from the die temperature sensor built into the TMP006. When this was attempted, it was found to work very poorly due to the low die temperature sensor noise, which gave rise to heavy ADC quantisation. In an effort to overcome the quantisation issue, an iterative fitting algorithm was developed in the GNU-Octave environment. Two assumptions were made; that a change in ADC value represented a point at which the true value was midway between the values either side of the change point, and that every ADC sample constrained the true value to within half of one least significant bit of that ADC value. The fitting process proceeded as follows; first the change points were used as fitting points for a cubic spline fit, then the resulting fit was checked against ± 0.5 least significant bit range limits. Any regions where the fit did not meet the constraints then had one more fit point added at their midpoint in time and at the maximum (i.e. ± 0.5LSB) possible deviation from the measured ADC value.

Once the new fit points had been added, the process was repeated until there were no regions where the fit deviated excessively from the measured ADC value.

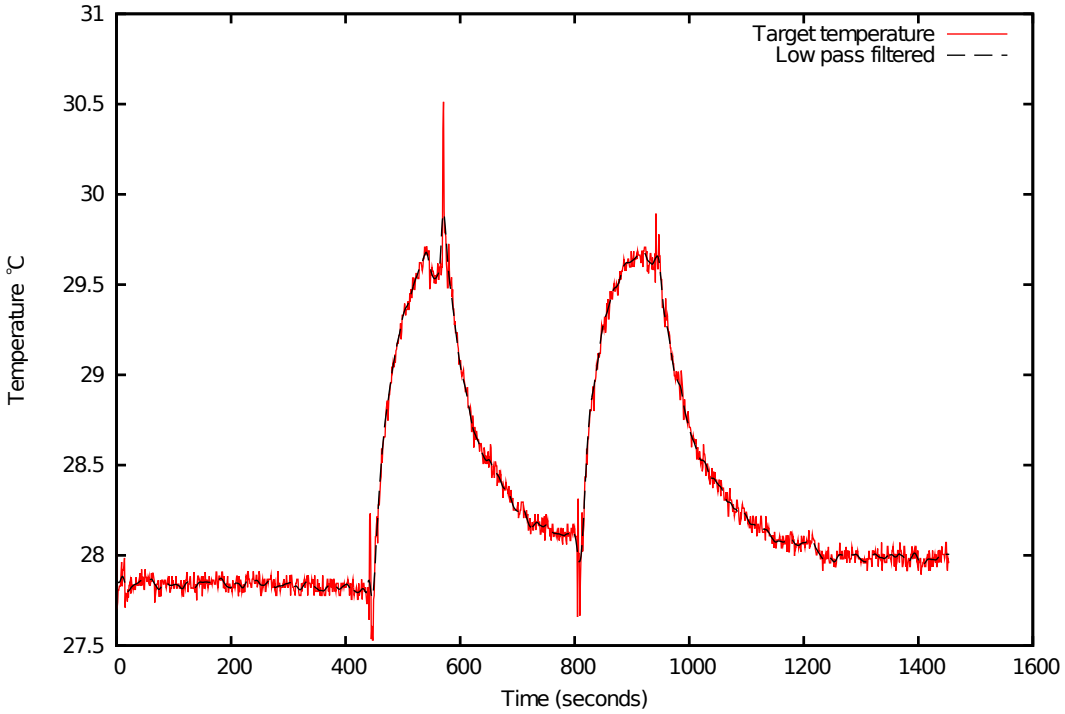


Figure 2.14: Filtered results from the non contact heater.

The TMP006 sensor was then heated with a hot air soldering tool to induce rapid temperature changes, with opaque plastic film placed over the sensor front to eliminate effects from long wave infra-red radiation. The interpolated target temperature and die temperature as a function of time were used to tune an adaptive correlation canceller. Once tuned, the cancelling network was applied to data recorded from the optical non contact heater.

It can be seen (figure 2.14) that the interpolation and correlation cancelling reduces both the amplitude and duration of the target temperature artefacts, and applying a low pass filter to the corrected data gives a particularly noticeable improvement.

Finally, the interpolated die temperature was subtracted from the low pass filtered data in figure 2.14 to give a measurement of optically induced heating. From figure 2.15 it can be seen that there is a small offset of approximately 1.25°C from the sensor, but that the optical heating is clearly detectable, the offset is stable, and the signal to noise is high enough for the exponential response to optically induced heating and cooling back to ambient to be clearly evident.

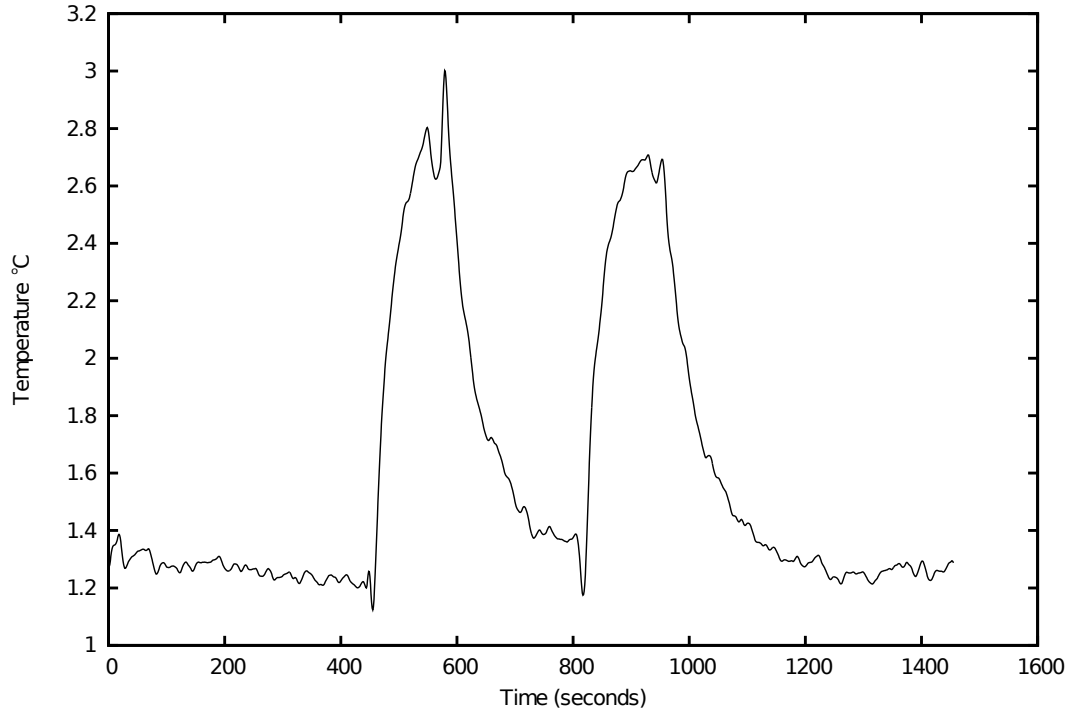


Figure 2.15: Filtered heating from the non contact heater.

2.6 Conclusion

Bioheat transport is a complex phenomenon, with multiple competing models in use to describe heat transport through perfused tissue. The modelling work in this chapter was limited to the simulation of thermal sensors using the Pennes equation, and so the validity of the results obtained could be called into question.

However, it has been shown in the simulations that a non invasive thermal disc sensor can provide useful information on volumetric perfusion, even in the face of inter patient and inter site variability in tissue thermal conductivity, and with imperfect adhesion of the sensor to the skin surface. The impact of changing environmental temperature and variability in core body temperature was not addressed in the simulation work. Neither was inhomogeneous perfusion or tissue thermal conductivity, such as that caused by a thick layer of epidermis. However, considering use on neonatal patients in an intensive care environment for example (one of the principal target subject groups), it should be noted that environmental temperature is relatively stable, and the skin is much thinner and more uniform than in adults.

It would appear from the number of published studies on the technique that thermal sensors are one of the most practical solutions to the problem of non-invasive perfusion monitoring. The fact that a heated thermistor technique is used in the Hemedex system [16], the only continuous perfusion monitoring device currently on the market, seemingly

demonstrates the practicality of thermal sensing. Numerous studies by Dittmar, Toumi et al using the "Hematron" sensor have shown strong correlation between a simple effective thermal conductivity measurement and laser Doppler PU (discussed in Chapter 6).

The laboratory experiments into optical heating demonstrate that it is a workable technique using commercial low cost sensors. However, the need for non contact heating was not demonstrated in the simulation work, so it may well be the case that optical heating is unnecessary and overcomplex, particularly when compared to exploration of methods for improved sensor adhesion, if adhesion and contact thermal resistance proved to be a problem.

It is clear that more research is required to assess the potential of thermal perfusion monitoring. Simulation to investigate effects of changing environmental and core body temperature would be a first step, but due to the complexity of bioheat modelling, volunteer studies would likely be required to progress much further. This would require prototype sensors to be constructed, using for example flexible copper coated poly-amide PCB material and thermistor heating would provide a route for rapid prototyping of sensors. It is notable that the "μHematron" device described by Toumi et al [23] uses a flexible PCB based sensor with soldered thermocouple elements. This construction approach avoids the need for well calibrated thermistors, as thermocouples provide a measurement of differential temperature that can be directly used with a differential temperature control scheme. Thermocouples using either soldered wires, or of a thin film construction more suited to mass manufacture may form a useful design element in a practical, production ready, thermal perfusion monitoring sensor.

Chapter 3

Use of Nuclear Magnetic Resonance for perfusion measurement

3.1 Introduction

By using spin labelling techniques, Nuclear Magnetic Resonance (NMR) may be suitable for the measurement of volumetric perfusion. This chapter starts by looking at the potential of a small Neodymium magnet based sensor for non invasive blood flow measurement.

A design study for a 10mm diameter permanent magnet NMR flowmeter is presented together with approximate (i.e. order of magnitude) estimates of performance. Some knowledge of NMR techniques and terminology is assumed.

3.1.1 Nuclear Magnetic Resonance and blood flow

As it consists primarily of H_2O, blood gives a strong ^1H signal of 42.6MHzT^{-1}. As the Larmor frequency is proportional to the field strength, a moving proton in a field gradient will have a varying Larmor frequency. This effect could be used to detect flow, for example if a 90° RF pulse is applied with a bandwidth b around a frequency f, then for a bandwidth sufficiently broad that T_1 and T_2 are insignificant over a time period of $1/b$, the frequency spectrum of the free induction decay will reveal flow as spectral components outside of the original 90° pulse bandwidth. This example is merely illustrative and several problems remain, for example diffusion may be very significant, and the signal to noise ratio and acquisition time need to be considered.

3.2 Instrument design

A simple and easily constructed design is a very high priority, meaning that a compromise with respect to field homogeneity and flux strength is necessary, resulting in less usable data. The most basic design is considered here (figure 3.1); a permanent magnet (B) placed on the skin surface with remote coil(s) to generate and sense RF.

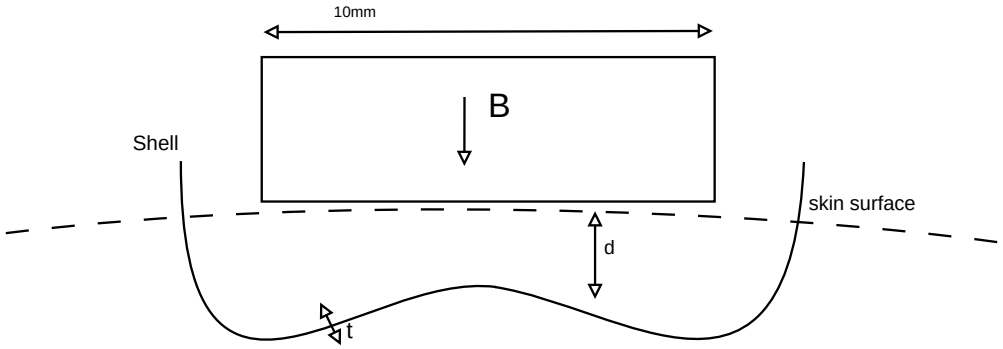


Figure 3.1: A simple permanent magnet NMR device

The magnetic field strength around a cylindrical permanent magnet decreases rapidly with distance, so by varying the frequency of a 90° RF pulse, it is possible to select thin volumes of tissue, of thickness,

$$t = \frac{b}{f_L \times G} \tag{3.1}$$

where f_L is the Larmor frequency of ^{1}H at 1T, b is the bandwidth of the RF pulse, and G is the magnetic field strength gradient. These volumes will be referred to as "shells" for the rest of this chapter. The depth d of a shell below the skin surface can be changed as a function of f the RF pulse centre frequency.

Assuming that the blood is able to reach thermal equilibrium (spin equilibrium) as it passes through the capillary bed, i.e. the transit time for blood flow through the capillaries is considerably greater than T_1 (of order 1500ms for blood), then blood flow through the tissue volume enclosed within the shell must exit passing through the shell, having ^{1}H proton spins that are in thermal equilibrium with the field. This would appear to be a fairly valid assumption, e.g. see [47]. However, if the flow velocity vector is approaching 90° to the local shell normal (i.e. the B field strength gradient vector), then the flow may still be unobservable even when the flow distance over the observation time is considerably more than the shell thickness. This is a source of error that appears hard to eliminate, although increasing the observation time and/or decreasing the shell thickness will help to resolve flow vectors lying closer to the normal. Using spin echo techniques, the observation time can be increased up to the T_2 limit, and the bandwidth can be decreased up to the limit where diffusion limits observations. Thus diffusion and

T_2 present a lower limit on observed flow speeds.

It will be shown that this limit is at or below the speeds typically encountered in human tissue, but in the extreme case of significant flow leaving the enclosed tissue volume via slow flowing capillaries orientated close to the field gradient normal, it appears impossible to avoid significant errors in the flow output.

Two further sources of error are differing sensitivity of the RF sense coil to different parts of the shells, and differing field strength from the RF exciter coil across the shells. If the sensor was to be used at two different sites without recalibration, this could present a source of error if there was any inhomogeneity in the distribution of blood vessels. Coils tailored to produce a homogeneous field/sensitivity over the sense region will help to mitigate these problems, but it would appear a compromise may have to be made with respect to device size.

Finally, a non constant field strength within the tissue volume enclosed by the detection shells will lead to a spatial variance in population inversion size. Thus in the case of blood flow to a region of higher field intensity, there will be a greater population inversion and hence a larger flow output. The ideal magnet configuration would have a region of homogeneous field surrounded by a region of field gradient allowing the selection of shells surrounding a volume of constant field strength where blood could reach thermal equilibrium.

In the first design study, a 10mm diameter N52 Nd-Fe-B will be considered, as this presents a good balance of mass, availability and price. The magnetic permittivity, μ_r, of N52 magnets is very close to 1, so externally induced flux can pass through the magnet without major distortions, but eddy currents will prevent high frequency RF from penetrating the magnet. This problem is not considered here, but one solution might be to use multiple smaller magnets in a polymer matrix. This would give the opportunity to experiment with magnet placement to increase field homogeneity [48] - for example using a CNC milled polymer enclosure to produce complex magnet alignments.

3.3 Pulse sequences

In an effort to devise an appropriate pulse sequence, the effect of tissue water diffusion needs to be considered. Diffusion will lead to movement of water molecules that could be misinterpreted as flow.

To a certain extent this could be treated as a DC offset on the measurements and removed, but this would be somewhat problematic, as diffusion is heavily influenced by body tissue structures, and thus may vary from patient to patient and as a function of

time. More significantly, due to the field gradient G, diffusion will lead to a dispersion of the distribution of Larmor frequencies. This will cause a T_2 effect - decreasing the amplitude of observed signal if significant diffusion into adjacent shells takes place.

Thus ideally a combination of the observation time over which the flow is detected (i.e. time between in initial 90° pulse and the observed free induction decay or spin echo), and thickness t of the observed shell is required which maximises the ratio of flow signal to diffusion signal.

Blood flow velocity in the capillary bed is typically of order $10^{-4}ms^{-1}$ or greater [49]. This can be compared to motion due to diffusion using Fick's second law (3.2)

$$\frac{\partial c}{\partial t} = D_M \nabla^2 c. \tag{3.2}$$

This relates the change in concentration c as a function of time t at any point to the Laplacian of the concentration. The constant D_M is equal to approximately 2.3x$10^{-9}m^2s^{-1}$ for water, and is found to take similar values in tissue [50], so long as the diffusion distance is less than the typical cell size.

It can be seen (from 3.1) that a narrowband RF pulse will select a thin shell, and the concentration gradient of 90° rotated spins will be high. From 3.2 this means that the time-scale of diffusion processes (T_d) scales with $(bandwidth)^2$. The thickness only scales with the first power of bandwidth, so the characteristic time for blood to move out of the selected shell (T_f) also scales linearly with ($bandwidth$). Thus at broad bandwidths flow will dominate over diffusion.

Unfortunately T_2^* scales with $bandwidth^{-1}$, so there may be problems using the basic pulse sequence described in the introduction. At this point it seems sensible to take a look at some real world numbers. K&J Magnetics publish data for N52 magnets [51]. Looking at product D64-N52, a gradient of 10^2Tm^{-1} and flux strength of 0.5T seem reasonable operating point assumptions. If we set $T_f = T_d$, then to a first order approximation,

$$T_f = \frac{\frac{b}{f_L G}}{10^{-4}} \tag{3.3}$$

$$T_b = \frac{(\frac{b}{f_L G})^2}{2.3 \times 10^{-9}}. \tag{3.4}$$

So, if $T_f = T_b$, then

$$b/(f_L G) = 2.3 \times 10^{-9} \times 10^4 = 2.3 \times 10^{-5} m. \tag{3.5}$$

Then, as $f_L = 42 \times 10^6 Hz$ and $G = 10^2$, $b = 96.6kHz$. Thus at least 100kHz of RF pulse bandwidth is needed to mitigate the effect of diffusion. Unfortunately this means that T_2^* is approximately $20\mu S$. T_1 for water is 4 seconds, so each shell must be left

for a period of order 4 seconds to return to thermal equilibrium (T_1 is usually 1500ms or less for blood, giving improved SNR). If shells of tissue are to be used sequentially, it can be seen that of order 1GHz of bandwidth must be used. This is not feasible, so there will be a reduction in signal to noise with increasing bandwidth as spins in our observed shells will be de-coherent most of the time, and hence producing no signal. Thus the bandwidth needs to be kept as low as possible without allowing diffusion effects to swamp the signal, so 100kHz was chosen. Note that in this analysis it was assumed that either the diffusion coefficient of the tissue water is much lower than that of pure liquid water, or that the diffusion distance of water in the tissue is limited by tissue structures. These assumptions are based on clinical data from diffusion tensor imaging.

Spin echo techniques exist as a solution to T_2^* signal decay. Unless mechanical means are used to rotate a permanent magnet, the field cannot be feasibly reversed for a gradient spin echo, so 180° RF pulse spin echo techniques must be used. Unfortunately for a cheap, compact instrument it would be difficult to achieve the necessary RF homogeneity to produce a chain of spin echoes, as the spins would experience differing rotations (i.e. not all are rotated 180°), so the spin axes become evenly distributed on the unit sphere after only a few spin echo sequences.

However, Spin echo does enable an increase in signal to noise ratio of several times, and also more importantly it allows an echo at relatively long periods of time after the initial 90° pulse, only limited by T_2 (around 180ms for venous and 255ms for arterial human blood).

This is important considering the condition on the time T_{ob} between initial shell selection and signal collection that needs to be satisfied in order for a signal to be distinguishable;

$$T_f \leq T_{ob}. \tag{3.6}$$

If this is not satisfied we will not observe the flow as the blood has simply not moved enough by the time we observe the signal. Solving equation 3.3 for $b = 5 \times 10^4$, then $T_{ob} = 230ms$. For this reason a sequence consisting of multiple 100kHz bandwidth 90° pulses followed 100ms later by a 180° RF pulse at the same frequency and with the same bandwidth will be considered, with interleaved RF pulses at differing frequencies spanning 5MHz of bandwidth from 15 to 20MHz, and with the sequence repeating at 0.25Hz to allow the spins in each shell to return to thermal equilibrium.

In figure 3.2, which plots the RF pulses and spin echoes in time frequency space, only 16 separate frequency bands are shown for clarity. The repetition time for the entire sequence would be of order 4 seconds, whilst the time between 90° RF pulses and spin echo at the same frequency would be of order 100ms. Flow would be detected as a

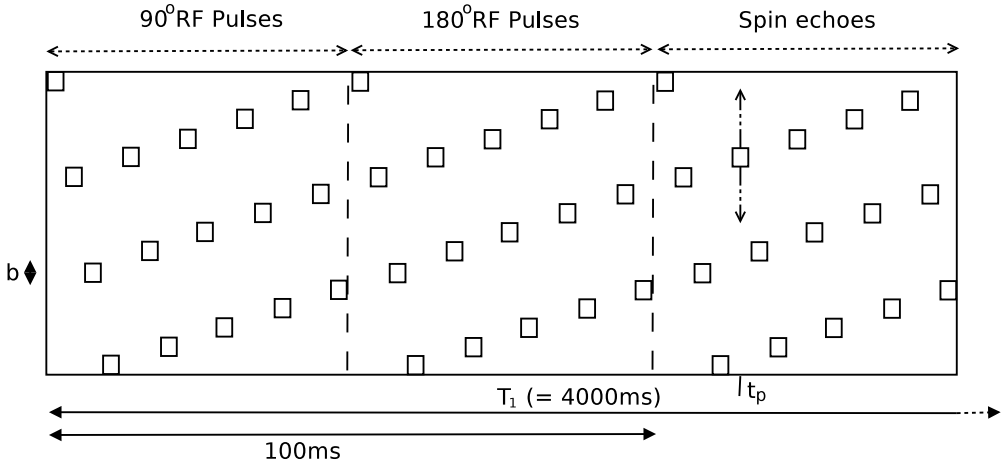


Figure 3.2: The proposed pulse sequence (T_1 is 4s for water, ∼1.5s for blood).

broadening of echo spectral components, this is shown by the dashed lines at probe time t_p.

3.4 RF coil driver requirements

As part of a feasibility check, it is obviously important to look at the RF driver requirements as well as the SNR margin for the receive coil. If a 180° pulse is to be used as part of a pulse sequence scheme, and occupies the same bandwidth as the 90° initial pulse, then the B field must have twice the magnitude, so the 180° is the strongest pulse. The B field from the RF exciter coil must be sufficient to cause a 180° rotation in 10μs, i.e. the Larmor frequency is at least $0.5/(10^{-5}) = 50kHz$. For ^{1}H, this corresponds to a B field of 1.2mT. Rounding this up to a 1.5mT peak field strength at coil centre, and assuming a 3 turn coil, then multiplying by a factor of 2 as there will only be a single coil (so the RF will be linearly rather than circularly polarized) gives 3mT.

$$B = \frac{\mu_0 NI}{2R} \tag{3.7}$$

solving for I,

$$I = \frac{6 \times 10^{-3} R}{3\mu_0}. \tag{3.8}$$

Assuming R=5mm,

$$I = \frac{4 \times 10^{-3} \times 5 \times 10^{-3}}{3\mu_0} = 5.3A. \tag{3.9}$$

The coil inductance will be [52]

$$L = N^2 r(\ln \frac{8r}{a} - 2 + Y)\mu_0, \tag{3.10}$$

where Y is a coefficient representing the degree of skin effect, a is the effective wire radius and r is the coil radius. Assuming Litz wire is used and the current is uniformly

distributed throughout the wire, then Y=0.25. To two significant figures (SF);

$$L = 5 \times 10^{-3}(\ln\left(\frac{40}{1}\right) - 2 + 0.25)\mu_0 N^2 = 8.7 \times 10^{-2} \times 4\pi 10^{-7} = 1.1 \times 10^{-7} H. \quad (3.11)$$

The peak voltage will be needed at 20MHz, the highest operating frequency. To two s.f.

$$V = I \parallel Z \parallel = I\omega L = 2 \times 10^7 \times 2\pi \times 8.0 \times 1.1 \times 10^{-7} = 110V, \quad (3.12)$$

so the peak voltage will be approximately 110V, with a current of 8A peak flowing through the coil.

The coil heating can be estimated by considering the resistive power losses. If the copper (resistivity 1.7×10^{-8}) Litz wire has a cross sectional area of $0.5mm^2$, then the total coil resistance for a 5mm radius with 3 turns will be $1.7 \times 10^{-8} \times 0.005 \times 3\frac{2\pi}{5\times 10^{-7}} = 4m\Omega$. So ohmic heating of the RF exciter coil will be around 260mW×duty factor, so around 0.2mW i.e. negligible. The peak voltage may be excessively high for convenient medical use, and the peak current would be hard to achieve, but it is interesting to note that the duty cycle is less than 0.1%. There seem to be two possible approaches to solving the coil current and voltage problem; reducing the flux gradient by around two orders of magnitude, or devising a pulse sequence that trades higher duty cycle for lower peak voltage and current.

3.5 Signal to noise ratio

The signal to noise ratio (SNR) was calculated assuming that thermal noise dominates over electromagnetic interference (EMI), as technologies exist for magnetically shielding the device in noisy environments, for example a magnetic shielding film such as Metglas 2714AZ could be used. A 100 turn copper coil of 0.1mm diameter wire was assumed. This has an effective series resistance (ESR) of approximately 6 ohms, giving a thermal noise floor of order $100pV(Hz)^{-1/2}$. This is very low, the noise floor of the best suitable RF amplifiers/ video opamps is around $2nV(Hz)^{-1/2}$. The number of turns could be increased and the wire diameter decreased to increase the SNR by better matching of the coil to the amplifier, but the number of turns was limited to 100 due to manufacturability considerations.

If the bandwidth is 100kHz, the total RMS noise in a single channel will then be $0.63\mu V$. This needs to be compared to the RMS voltage generated by the spin precession in the probed volume. If the field strength is 0.5T, δE between spin up and down will be 1.3×10^{-26}J. The Boltzmann distribution states that the ratio of spin up to spin down is equal to

$$\frac{N_{up}}{N_{down}} = e^{\frac{\delta E}{k_B T}}. \quad (3.13)$$

Solving this for T=310K, to 2 s.f.

$$\frac{N_{up}}{N_{down}} = 1 + 3.2 \times 10^{-6}. \tag{3.14}$$

So there is a fractional population inversion of 1.6×10^{-6} (i.e. 1 spin in 1.6×10^{-6}). The number of spins can be found from the shell thickness, equation (3.1). For a 100kHz bandwidth and G=100, this gives 24μm. The approximate volume will thus be $\pi(\frac{1}{2})^2 \times 24 \times 10^{-4} ml = 1.9 \times 10^{-3} ml$. If the assumption is made that the tissue has such a high water content that it can be approximated as purely water, then there are $\frac{N_A}{18} 1.9 \times 10^{-3} = 6.3 \times 10^{19}$ H_2O molecules, or 1.26×10^{20} ^{1}H nuclei. Thus (to 2s.f.) we can manipulate a population of $1.6 \times 10^{-6} \times 1.3 \times 10^{20} = 2.0 \times 10^{14}$ nuclei per shell.

As the magnetic moment of a ^{1}H nuclei is $1.4 \times 10^{-26} JT^{-1}$, the total manipulated magnetic moment is;

$$1.4 \times 10^{-26} \times 2 \times 10^{14} = 2.8 \times 10^{-12} JT^{-1}. \tag{3.15}$$

So in the case of a coil with 100 turns (N), and introducing a flux linkage factor $l = 0.1$ as an estimate for the fraction of flux from an average spin that could be linked through the sense coil, the induced voltage at 17.5MHz will be (to 2SF);

$$2.8 \times 10^{-12} \times 1.75 \times 10^{7} \times 2\pi N l \times \mu_0 = 1.2 \times \pi 10^{-9} = 3.8 \times 10^{-9}. \tag{3.16}$$

So the signal to noise ratio per shell will be:

$$\frac{3.8 \times 10^{-9}}{6.3 \times 10^{-7}} = 6.0 \times 10^{-3}. \tag{3.17}$$

However, there are 50 probed shells at different frequency bands, spanning a 5MHz total bandwidth. Thus the signal to noise ratio in 3.17 can be multiplied by a factor of $\sqrt{50} = 7.1$ giving;

$$6.0 \times 10^{-3} \times \sqrt{Shells} = 4.3 \times 10^{-2}. \tag{3.18}$$

The complete sequence of interleaved spin echoes will repeat every 4 seconds or thereabouts (this may be reduced to 1.5 seconds for blood), allowing a fraction $1 - 1/e$ of the spins to decay to thermal equilibrium over the T_1 time period. Thus the realizable SNR for a continuous monitoring process will be even lower at around

$$4.3 \times 10^{-2} \times \frac{e-1}{e} = 2.8 \times 10^{-2}. \tag{3.19}$$

Assuming perfusion values towards the low end of those typically seen in the hand; around $10ml/100g/minute$, and using the earlier approximation whereby non aqueous tissue components are ignored, the SNR will be:

$$2.8 \times 10^{-2} \times \frac{10}{100 \times 60} = 5.1 \times 10^{-5}. \tag{3.20}$$

thus the perfusion signal at a bandwidth of 0.125Hz will be massively masked by the noise, and it is clear that such a simple device is not capable of producing medically useful data.

3.6 High Homogeneity magnet configuration

High homogeneity configurations would allow the bandwidth of each "shell" to be decreased, thus reducing the demodulated thermal noise. To investigate the scaling of SNR with field homogeneity, the relationship between diffusion and perfusion needs to be explored in more detail.

There has been considerable research into permanent magnet configurations with remote homogeneous fields, Marble et al [48] and Manz et al [53] discuss possible configurations. However, most existing research appears to have concentrated on creating a region of very high homogeneity, for NMR chemical analysis. These requirements differ slightly from those for a flow meter, where the overall homogeneity is less important than a large region of low gradient. It would also be advantageous to have the field orientated in the same plane as the skin surface, so as to allow compact coils to be used. For this reason the design discussed in Perlo et al [54] seems a good choice if the design could be altered to enlarge the homogeneous region at the expense of homogeneity. Blank et al [55] discuss an NMR system built into a catheter, the use of a varactor for variable coil tuning is of note, as is the extremely high field gradient.

Considering the earlier results, where it was found necessary to use spin echo due to extremely short T_2^*, it can be seen that spin echo will still be necessary even with a dramatic increase in coherence due to improved homogeneity. Thus it will be assumed that the pulse sequence considered earlier will remain unchanged, and only the bandwidth will be altered.

The T_1 time will not be influenced by changes in the applied field, and thus it will be assumed that a single spin echo sequence with the same time to echo will be used. Multiple spin echoes using pulse trains are feasible with high homogeneity systems, but these require a high homogeneity field from the RF coil in order to maintain coherence between the selected spins and so generate an echo after multiple 90° or 180° pulses. The thermal diffusion will likewise be unaffected by field homogeneity. Thus it can be seen that the physical width of a "shell" will not change, only the frequency bandwidth occupied by the spins within it, and by implication the signal to noise ratio.

If the noise is approximated as thermal noise with a flat frequency spectrum, it then follows that the signal to noise ratio will improve with the square root of the improvement in field homogeneity. From the signal to noise ratio this would imply an improvement in homogeneity by a factor of over 3×10^8 or a value of order $3 \times 10^{-7} Tm^{-1}$. This is impractical with small permanent magnet configurations, but fields of this order of homogeneity are commonly generated in large scale whole body NMR magnets (the derivation of this homogeneity value also serves as a simple demonstration of the feasibility of diffusion tensor imaging with state of the art whole body MRI).

It is clear that improvements in field homogeneity alone will not be sufficient to obtain useful perfusion information with a small permanent magnet NMR device. One option that remains open for increasing signal to noise is to increase the probed tissue volume. This inevitably leads to an increase in sensor size, but if the probed volume can be increased, the number of observable hydrogen nuclear spins will increase in proportion, and so a larger signal to noise ratio will be achievable. A simple technique for analysis of the volume scaling might be to model a linear scaling of the sense volume along all three of its dimensions. For a linear scaling factor of L along all axes, and keeping the field homogeneity unchanged, it can be seen that the number of atomic spins in a "shell" scales with L^2, and the number of probed "shells" scales with L, giving a signal to noise ratio scaling of $L^2 \times L^{0.5} = L^{2.5}$. If the earlier sensor was scaled by a factor of ten, to a 5cm radius sense area, then this simple scaling law would suggest that a useful signal to noise ratio could be reached without the necessity for an extremely homogeneous field.

A more thorough analysis would examine the number of "shells". Published studies on permanent magnet NMR have reached gradients of order $0.01 Tm^{-1}$ [56]. In our model this would give a bandwidth of order 10Hz per shell, allowing many shells to be used without the RF coil requiring an excessively large bandwidth, and the dimension of the sense volume along the field gradient would only be limited by the size of the homogeneous field volume created by the sensor. Looking again at the signal to noise ratio of 5.1×10^{-5} derived earlier, if the sensor was scaled up to give a sense volume of 5cm diameter (i.e. a diameter scaling of 5), the thickness of a shell was unchanged at $24 \mu m$, and the number of shells was set to 10^4, giving a 100kHz total bandwidth, then the signal to noise ratio would be improved by a factor of order $5^2 \times \sqrt{\frac{10^4}{50}}$ from the increased sense volume. The reduction of bandwidth from 100kHz to 10Hz for each shell would give an additional improvement of a factor of $\sqrt{\frac{10^5}{10}} = 10^2$.

Putting this all together to scale our previous signal to noise ratio results in the following SNR:

$$5.1 \times 10^{-5} \times 10^2 \times 5^2 \times \sqrt{\frac{10^4}{50}} = 1.8. \tag{3.21}$$

This is a medically useful value, especially if only long term perfusion trend data are desired. It may also be possible to use T_2 effects to distinguish different blood flows due to the relatively long delay before the spin echo, which could be made to occur approximately midway between the T_2 for oxygenated and de-oxygenated blood.

US Patent 4,613,818 [57] describes a very similar device to this, and papers by Battocletti et al [58, 59] describe similar whole limb based NMR perfusion monitoring devices. The device described in [58] differs somewhat from the concepts described here in that the "magnetization" (i.e. the creation of a population inversion) and the sensing of the spins take place at physically separate locations.

Some authors have proposed use of Earth field NMR in conjunction with a physically separated permanent magnet used solely to create a population inversion, for example Appelt et al [60]. Based on this idea, a simple permanent magnet attached to the skin surface at the region of interest could be used to polarize perfusing blood. A large rotation and sensing coil system, perhaps encompassing the entire body could then be used to detect the H_1 nuclear spin, giving a signal intensity proportional to the blood flow through the field affected tissue, and thus a measure of local perfusion. Unfortunately such a system would be highly impractical for use in a hospital environment, since the earth field strength and direction inside a hospital building would inevitably be highly variable, and it would be extremely difficult to shield the open coils from stray radio frequency emissions from equipment.

A higher homogeneity magnet configuration would allow the bandwidth to decrease and hence the "driver" coil pulse duration to increase, giving lower current and voltage in the coil. Previously 110V and 8A were found as approximate figures for a 100Tm^{-1} field gradient, so if the field gradient could be decreased down to around 5Tm^{-1} or lower, then from equation 3.1, the required bandwidth will decrease down to 5kHz. The 180° pulse current will then be 400mA for a 3 turn coil, and peak voltage will be 5.5V. These figures are clearly more reasonable for a medical device.

3.7 Conclusion

A design study for a simple 10mm diameter permanent magnet based NMR sensor for the purpose of non invasive blood flow monitoring in human tissue indicates that no useful data can be obtained, as the signal would be massively swamped by noise.

Non-conventional schemes involving physically separated polarization and sensing have been employed in experimental perfusion sensing devices, but such a scheme seems unlikely to be practical in a hospital environment due to stray magnetic fields and electromagnetic interference.

High homogeneity schemes would appear to be practical based on the low field homogeneities reached in published studies. However, the probed tissue volume would have to be large in order to reach a useful signal strength. In practice this would probably limit the use of such a device to spot checks of perfusion to an entire limb, and so continuous monitoring would be quite impractical. US patent 4,613,818 describes such a device.

The problem of excessively high drive voltage with the basic dipole magnet approach means that it would probably prove necessary to use multiple magnets arranged in a polymer or high magnetic permeability matrix, but a literature review indicates that generation of a high homogeneity field is feasible with similar schemes. For example if the design is based on that in [54], then "Kool-mμ" - see [61] would seem a possible material for backing plate, having a high saturation flux of 1.1T and being non-conductive.

From the low cost of suitable Neodymium magnets, it would appear that the sensor unit would be comparable in manufacturing cost to pulse oximeters and similar optical devices. Given a suitably low field gradient of around $5\mathrm{Tm}^{-1}$, no serious risks are apparent, as currents and voltages are all within safe limits (the risk of overheating the exciter coil in the event of a driver circuit failure could be mitigated by redundant fuses and/or a small series resistor).

Placement of the RF coils is complicated by the requirement of a homogeneous field and sensitivity across the sense volume. This requirement will probably cause the device diameter to increase considerably beyond the 10mm initially discussed.

It would thus seem reasonable to conclude the application of NMR for perfusion assessment has only been proven in high field MRI scanners. Any new device using for example permanent magnets would represent an extremely high risk research investment, and it would appear that even if a sensor could be successfully developed, the most likely outcome would be a considerably larger sensor than would be optimal for peripheral monitoring, with high susceptibility to electromagnetic interference from the surrounding environment.

Chapter 4

Trans-cutaneous blood gas monitoring for perfusion sensing

4.1 Introduction

Trans-cutaneous blood gas sensing attempts to measure the partial pressure of blood gasses non-invasively through the skin surface. Its use has been explored both as a direct indicator of volumetric perfusion, and as a method of assessing perfusion adequacy to the tissue. Trans-cutaneous oxygen and carbon dioxide sensors are available as off-the-shelf medical devices, and numerous studies have demonstrated their potential for monitoring of local perfusion adequacy. However, in current clinical practice they are only routinely used for the monitoring of arterial blood, providing a metric of whole body metabolic function.

Trans-cutaneous sensing of biologically inert tracer gasses such as hydrogen has been attempted in several studies, for example Varghese et al [62] with a solid state H_2 sensor, and Harrison et al [7] used a modified PtO_2 sensor to monitor H_2 partial pressure. Here a measure of absolute volumetric perfusion can be attempted through use of a washout technique. Washout or clearance of a dissolved tracer gas has been employed in numerous studies to measure perfusion via detection of an inhaled radioactive gas such as Xenon that dissolves in the arterial blood and diffuses into the tissue, with trans-cutaneous detection of the gamma rays emitted from the Xenon decay providing an indication of the tissue Xenon concentration as the gas is flushed out via gas free perfusing blood once the patient switches to inhalation of tracer free air. In similar fashion, a liquid solution of radioactive Iodine 125 can be injected into the arterial blood supply and employed as a tracer. Tracer washout is regarded by many as a gold standard perfusion metric, for example Harrison et al [7] described Xenon clearance one of the few established methods

capable of measuring volumetric perfusion, and employed ^{125}I 4-iodoantipyrine clearance (IAP) as one of the benchmarks for their Hydrogen clearance measurements. Similar clearance techniques are also feasible by using ultrasonically tracked micro-bubbles or florescent die (such as Indocyanine green) injected into the arterial blood supply and traced via ultrasound or florescence spectroscopy respectively. The majority of these tracer techniques are clearly either invasive, requiring injections into the arterial supply to the probed tissue, or dangerous, requiring use of radioactive isotopes and so inherent radiation exposure to the patient. Thus it is clear that with the exception of hydrogen clearance, these methods are not suitable for use in continuous non-invasive perfusion monitoring, but may still serve as a benchmark for any new candidate device.

It should also be noted that all published inert tracer studies have used gas inhalation via a respirator mask followed by trans-cutaneous gas detection, and evidence for successful trans-cutaneous tracer gas diffusion into the skin could not be found.

4.1.1 Current clinical practice

Trans-cutaneous blood gas sensing is a widely used technique in neo-natal intensive care [63].

Although trans-cutaneous partial pressure of oxygen (PtO_2) was formerly in widespread use as a vital sign metric, in recent years it has been widely superseded for many routine applications by pulse oximeters based on differential absorption spectroscopy. However, PtO_2 is still in wide use for continuous monitoring of oxygen partial pressure, as pulse oximetry can only measure the oxygenation state of haemoglobin, and so cannot reliably detect hyperoxic events, where the haemoglobin is fully oxygenated and the partial pressure of oxygen rises to dangerous levels. This risk is especially significant for neonates receiving assisted ventilation (e.g. CPAP - continuous positive airway pressure), who may be being ventilated with an elevated oxygen level (Dear et al [64]).

Partial pressure of carbon dioxide ($PtCO_2$) is another blood gas metric that is widely monitored in current practice, though not as commonly as PtO_2. $PtCO_2$ is typically used only as an indicator for hypoventilation in clinical care. However, there is evidence to suggest that elevated tissue CO_2 is also an indirect indicator of poor perfusion, for example Brunstler et al [65] found that ($PtCO_2$) was significantly lower than invasively measured arterial $PaCO_2$, with $PtCO_2$ being "grossly" elevated above $PaCO_2$ in the case of systolic blood pressure below 15mmHg in a group of 24 sick neonates, and with the arterial to trans-cutaneous drop correlating strongly with reduced sensor heating (and by implication reduced perfusion, due to the response of the skin to sensor temperatures above 40°C). This correlation might be expected from basic theory, as with a reduced

perfusion tissue metabolism will give rise to a larger partial pressure of carbon dioxide in the venous blood.

In modern critical care $PtCO_2$ is typically measured using a combination $PtCO_2$ and PtO_2 sensor, as this means that only one site has to be exposed to the sensor, which is typically heated to a temperature of between 42°C and 44°C and thus poses a risk of skin damage. Lubbers [66] provides a detailed explanation of the rationale behind heating of the skin. For PtO_2 measurement heating serves to both increase the permeability of the epidermis to oxygen, and dramatically increase perfusion by inducing vasco-dilation. This means the tissue oxygen consumption is negligible relative to the arterial oxygen supply, and thus tissue O_2 partial pressure approximates that of the arteries. For $PtCO_2$ measurement, heating serves both to stabilise the sensor, and to increase epidermal permeability. Some authors have explored unheated $PtCO_2$ sensors (e.g. Cabal et al [67]), with some apparent success, particularly on neonates where the epidermis is thinner and so more permeable to CO_2. Despite its potential to avoid the risk of permanent tissue damage from heating, unheated $PtCO_2$ monitoring has not seen widespread use, apparently due to the prevalence of combination $PtCO_2$ and PtO_2 sensors, with the O_2 sensor requiring tissue heating as shown by Cabal et al.

4.2 Theory of operation

4.2.1 The Clark electrode

Although alternative solid state electrodes have been demonstrated in the laboratory (for example Lanthanum flouride based by Yamazoe et al [68]), the Clark electrode (Clark et al, [69]) is the standard trans-cutaneous oxygen sensor. This is actually a cell arrangement, rather than a single electrode. It consists of a Ag/AgCl reference anode, and a catalytic platinum cathode consisting of platinized platinum; platinum foil coated with a thin layer of highly porous platinum, usually via electroplating techniques. An electrolyte solution, typically KCl dissolved in water is used between the two electrodes, and a thin, gas permeable membrane usually comprised of polythene or Teflon (PTFE) is placed between the cell and the patients skin surface. Contact gel is typically used to exclude air from the membrane surface, and in modern use the cell itself is usually a thin film design using a gelled electrolyte, the disposable electrodes plugging into a reusable back cover containing a heater, temperature sensor, and cable to the bedside monitor. Disposable adhesive attachment rings are then used to secure the sensor unit to the patients skin. Due to the potential for long term skin injury from prolonged skin heating, common practice is to move the sensor site every 3 to 4 hours, reattaching

the sensor with a fresh attachment ring. The thin film cells are also regularly replaced, sometimes each time the sensor location is moved.

The cell functions by catalytic reduction of oxygen at the cathode, and conversion of Ag to AgCl at the anode.

Anode:
$$O_2 + 2H_2O + 4e^- \rightarrow 4HO^-, \tag{4.1}$$

cathode:
$$4Ag + 4Cl^- \rightarrow 4AgCl + 4e^-, \tag{4.2}$$

overall reaction:
$$O_2 + 2H_2O + 4Ag + 4KCl \rightarrow 4AgCl + 4KOH. \tag{4.3}$$

Typically a fixed potential is applied to the cell, with the cell current then proportional to oxygen availability at the platinum surface, which can be calibrated against partial pressure of oxygen to give an accurate measurement. Voltages between 400mV and 1200mV are usable, but 800mV is typically used, giving optimum sensitivity to O_2 versus other species, sensor drift, and temperature effects. This fixed potential operating mode is termed polarographic or alternatively amperometric, as electrode current serves as the sensor output.

As can be seen from equation 4.3, in operation the sensor electrolyte becomes increasingly alkaline, and Cl^- ion availability decreases. This biases the sensor current, and is one of the main reasons for the limited operational lifetime of the sensor. The sensor lifetime can be increased by stabilising pH with a phosphate buffer [70], but the Ag/AgCl anode is still prone to being covered with a solid AgCl film, due to the extremely low solubility of AgCl in water.

Galvanic cells such as the Hoeffer cell and the Mackereth cell [71] are often used to measure dissolved oxygen in laboratory settings. These cells operate on the oxidation of a Zn or Pb cathode, and have superior long term stability to the Clark cell. ZnO and PbO form solid precipitates that do not cling the cathode or dissolve significantly in the electrolyte. Thus cell lifetime is limited only by the supply of metal. Unfortunately the Macreth cell uses concentrated NaOH electrolyte and a Pb electrode, so is too hazardous for use in trans-cutaneous sensors. The Hoeffer cell uses a Zn electrode and NaCl electrolyte, so it safe for use in biomedical sensing. Unfortunately CO_2 dissolved in the electrolyte causes biasing of the sensor [70], and for this reason Clark cells are used in almost all current commercial oxygen sensors.

4.2.2 Trans-cutaneous CO_2 sensors

Trans-cutaneous CO_2 ($PtCO_2$) sensors operate on a pH measurement principle. After diffusing through a gas permeable membrane (thin polythene or Teflon), CO_2 dissolves into an aqueous bicarbonate buffer solution, changing the equilibrium between carbonic acid (H_2CO_3) and CO_2/water [72]

$$CO_2 + H_2O \rightleftharpoons H_2CO_3 \rightleftharpoons HCO_3^- + H^+. \tag{4.4}$$

A pH sensing glass electrode can then be used to monitor the pH of the solution, and thus calculate CO_2 partial pressure. Measurement cells for pH consist of two identical electrodes, typically Ag/AgCl, with one in contact with the reference fluid via a salt bridge, and the other via a thin glass membrane mobile to Na^+ ions. Exchange of Na^+ ions with H^+ at the glass surface means the glass is H^+ selective, and thus the cell output voltage serves as a measurement of H^+ concentration in the fluid, and thus allows it's pH to be determined.

In modern CO_2 sensors the unreliable and expensive glass electrode and electrolyte system has been replaced by solid state pH sensing. These solid state sensors are mostly metal/metal oxide electrode systems with a cation selective polymer coating. Fog et al published an overview of solid state pH sensing using such systems in 1984 [73], and such a pH sensor is described in US Patent 5,218,304 [74].

4.2.3 Use of the Clark electrode for H_2 sensing

By applying a reverse biasing voltage to the previously described Clark electrode configuration, it is possible to change the cell chemistry and detect H_2, this is commonly referred to as a "Clark hydrogen electrode" in relevant literature. The necessary changes to adapt a conventional PtO_2 probe are described by Harrison et al [7]. With a reverse biasing voltage applied to the cell, the reactions are now;

anode:
$$2Ag^+ + 2e^- \rightarrow 2Ag, \tag{4.5}$$

cathode:
$$H_2 \rightarrow 2H^+ + 2e^-, \tag{4.6}$$

overall reaction:
$$H_2 + 2AgCl \rightarrow 2Ag + 2HCl. \tag{4.7}$$

For H_2 detection a biasing of approximately 600mV is applied between the platinum cathode and silver anode. The overall reaction 4.7 increases the acidity of the electrolyte, which limits the lifetime of the sensor in a similar fashion to the Clark oxygen electrode.

4.3 Trans-cutaneous blood gas and perfusion

There are numerous means by which trans-cutaneous gaseous exchange and sensing may be of use for perfusion monitoring:

4.3.1 H_2 washout techniques

Perhaps the most conceptually simple technique is the hydrogen clearance method, this was first developed by Aukland et al [75]. Pasztor et al demonstrated its use for invasive cerebral perfusion sensing [76], and many more recent studies on CBF have employed H_2 clearance. More recently, Harrison et al developed a non invasive clearance technique [7]. In Harrison's study, patients inhaled a mixture of air with 2% H_2 gas (below the flammability limit) for 10 or 20 minutes via a respirator. This brought their body tissues into equilibrium partial pressure of hydrogen with the inhaled gas via diffusion exchange between the tissue and capillary bed with perfusing blood. After the gas inhalation phase was complete, the patients were allowed to return to breathing room air. The tissue H_2 partial pressure was then measured trans-cutaneously using a conventional oxygen sensor modified to increase anode reactivity with hydrogen, and driven with reverse bias relative to the O_2 sensing configuration. The current through the cell was found to be proportional to H_2 partial pressure in calibration tests, so a simple linear scaling was used to convert the measured current to H_2 partial pressure.

Experimental data showed a clear exponential decay, as would be expected as hydrogen was "flushed out" of the tissue by perfusing blood (similar to the thermal "flush out" term in the Pennes bioheat equation). Some Hydrogen is also lost due to diffusion out through the skin to the room air, but this effect was found to be negligible due to the large size of the sensor relative to the diffusion length causing it to block gas diffusion out of the skin. Application of an arm cuff to restrict blood flow confirmed this conclusion, as the $PtcH_2$ decay ceased when the cuff was inflated. The volumetric perfusion derived from exponential fitting to the current decay was found to correlate strongly with perfusion measured via an alternative invasive technique of injecting micro-bubbles and monitoring their path using ultrasound. However, the trans-cutaneous H_2 derived perfusion was approximately three times larger. This increase was speculated to be due to the perfusion increase induced by heating the skin to 40°C. Operation of the sensor at body temperature was found to be unreliable. The precise mechanisms by which a sensor temperature increase from 37°C to 40°C apparently caused such a significant increase in skin permeability is unclear, but the dramatic behaviour change over such a small temperature change was speculated by Harrison et al to apparently

indicate a physiological origin, rather than for example a simple physical properties based increase in epidermal permeability with temperature.

One potential explanation is that as the temperature is elevated above core body temperature, sweat glands and pores dilate, allowing much easier gas exchange through the skin i.e. enhanced diffusion. For example Singh et al [77] discuss trans-follicular routes as an important pathway for trans-dermal drug delivery.

Washout techniques, using tracers such radioactive gasses or liquids, fluorescent dye, or ultrasonically traceable micro-bubbles injected into a central artery form the clinically accepted "gold standard" for perfusion assessment [78]. The fact that hydrogen is relatively biologically inert makes it perhaps the most attractive perfusion tracer. H_2 washout might even be feasible for continuous monitoring; for use on a patient already undergoing assisted ventilation, since it would be feasible to introduce a low concentration of H_2 gas into their air supply. Unfortunately this would introduce flammability and explosion hazards in the event of an equipment malfunction. There is also the requirement of a heated sensor, with its inherent risk of tissue damage by overheating.

It may be notable that the research by Harrison et al was conducted with adult patients. There is evidence to suggest that neonatal skin has higher permeability than that found in adults, as might be expected from this thinner epidermis. For instance Cabal et al found unheated $PtcCO_2$ sensing to be feasible in neonates [67]. There has also been much research into methods to increase skin permeability, particularly with a view to use in trans-dermal drug delivery systems, for example Moser et al [79] examined the effect of a range of gels applied to the skin surface on permeability. The extent to which similar techniques are already in use in commercial trans-cutaneous sensors is unclear, as the contact gels are proprietary, but the existence of patents such as US Patent 4,488,557 [80] suggests this may be the case.

One possible route to a more practical trans-cutaneous H_2 clearance sensor would be to generate small quantities of hydrogen electrolytically within the sensor, a method first used invasively in 1974 by Stosseck et al [81]. This would result in a H_2 clearance sensor with many similarities to a simple thermal sensor measuring input heat flux at a constantly maintained temperature elevation. The principal differences would be that unlike the baseline temperature, the H_2 baseline level is zero, so there would be no requirement for the "ambient" tissue hydrogen level to be measured. The diffusion coefficient for hydrogen in tissue is an order of magnitude lower than the thermal diffusion coefficient ($5 \times 10^{-9} m^2 s^{-1}$ for H_2 as opposed to thermal diffusion of around $5 \times 10^{-8} m^2 s^{-1}$ or slightly higher).

For trans-cutaneous sensing this means that in the unperfused state the loss of hydrogen

will be relatively insignificant, in contrast to thermal monitoring with a heated sensor, where heat flux due to tissue thermal conductivity must be calibrated out. Unfortunately the low diffusion coefficient also means hydrogen will not penetrate as deeply into the skin as heat flux, meaning a hydrogen based sensor would be considerably more strongly influenced by skin properties and coupling problems than a thermal sensor, and would be sensitive only to perfusion in the 1mm or so of tissue below the skin surface. It should however be noted that due to the very low flux of hydrogen through the skin in an unperfused state, any uncertainties in skin permeability would manifest as sensor gain errors, rather than as a mixture of offset and gain error as is the case with thermal. The high diffusion coefficient of hydrogen in air, $6 \times 10^{-4} m^2 s^{-1}$, is some five orders of magnitude greater than that of hydrogen in typical tissue, whilst hydrogen molar solubility at $2kPa$ partial pressure and 37°C temperature is roughly equal. This compares to a two order of magnitude difference in thermal conductivities between air and tissue, meaning that any air gap under the sensor would have an overwhelmingly larger influence on sensor data from a trans-cutaneous hydrogen sensor than from a thermal sensor. Nevertheless, problems of air gaps would seem to have been mitigated relatively effectively in trans-cutaneous O_2 and CO_2 sensing.

4.3.2 Trans-cutaneous CO_2 as a perfusion indicator

The blood perfusing through the capillary bed serves three principal metabolic roles;

1. Supply of oxygen to the tissue (via haemoglobin)
2. Removal of carbon dioxide from the tissue (via bicarbonate buffer system)
3. Supply of glucose to the tissue (blood glucose) and other nutrients/electrolytes.

The first of these roles is of course a subject of great interest, with devices such as for example pulse oximeters and trans-cutaneous oxygen sensors used to record arterial oxygen in real time. The last role is encountered in conditions such as diabetes, but the second role; that of CO_2 removal, is perhaps the least discussed of these three metabolic blood supply roles in the context of specific medical conditions.

Just as low perfusion causes tissue oxygenation (i.e. the ratio of HbO_2 to total Hb) to decrease, as tissue metabolism reaches a new equilibrium oxygen tension, a higher equilibrium carbon dioxide tension will be reached in similar fashion due to the reduced ability of the perfusing blood to remove carbon dioxide. In clinical science two distinct resultant conditions are recognised: respiratory acidosis; pH decrease caused directly by insufficient CO_2 removal, and metabolic acidosis; induced by lack of oxygen supply

relative to the metabolic rate, a condition that is commonly encountered in the muscle tissue during vigorous exercise. In the second case the pH shift is caused by changes in the metabolic pathway [82], but CO_2 partial pressure is still shifted due to H^+ uptake by the bicarbonate buffer system [83, 84].

Cabal et al [67] correlated trans-cutaneous CO_2 with end tidal CO_2, finding that although there was much variability in both measurements of CO_2, the discrepancy between the two was strongly correlated with hypoxaemia at the trans-cutaneous measurement site. Tatevossian et al [8] found correlation between low trans-cutaneous O_2 and high CO_2 and negative outcomes for intensive care patients. Tremper et al [85] conducted a similar study, including patients with shock and postoperative patients. In a 1980 study, the same authors demonstrated trans-cutaneous CO_2 increase during induced shock in dogs [86]. Eberhard et al [87] found strong correlation between arterial partial pressure of CO_2 (measured invasively), and trans-cutaneous CO_2, the trans-cutaneous pressure being significantly higher, which was hypothesised to be due to the metabolism of tissue local to the trans-cutaneous sensor. This might be expected as tissue metabolism gives rise to a larger partial pressure of carbon dioxide in the venous blood than the arterial, and $PtcCO_2$ will approximate a mixture of venous and arterial blood.

4.3.3 Trans-cutaneous O_2 as a perfusion indicator

Trans-cutaneous O_2 is typically used as a measure of arterial partial pressure of oxygen. The use of sensor heating serving both to increase skin permeability and to induce vasodilation, increasing perfusion beneath the sensor, this "arterialises" the tissue, bringing tissue partial pressure of oxygen close to arterial partial pressure.

However, numerous authors have investigated the difference between true arterial partial pressure of oxygen (PaO_2) and the trans-cutaneously measured tissue value. As this difference should reflect that between the metabolic uptake of oxygen and the arterial supply, a lower trans-cutaneous tension, or more specifically a greater difference between the arterial and trans-cutaneous tension might be expected to reveal low perfusion. White et al [88] used dual site $PtcO_2$ sensing to identify peripheral vascular disease, and Tatevossian et al [8] found low $PtcO_2$ correlated with poor patient outcomes, and speculated that $PtcO_2$ was an indicator of tissue hypoxia, although their study did not measure baseline arterial oxygen tensions. Gottrup et a [89] found an inverse correlation between trans-cutaneous tissue oxygen partial pressure and perfusion in a study involving simultaneous monitoring of arterial and trans-cutaneous O_2 tensions in dogs.

The excess perfusion induced via heating is useful in conventional trans-cutaneous de-

vices as it brings the partial pressure closer to the arterial pressure. Provided there was an independent means of assessing arterial oxygen, partial pressure data that more closely approximated tissue baseline partial pressure would be much more useful for perfusion estimation. For example, a conventional dual wavelength pulse oximeter could be combined with a trans-cutaneous oxygen sensor, and oximetry derived Hb saturation could be converted to an arterial PpO_2 estimate via a Hb oxygenation curve (although this would be somewhat prone to error due to e.g. unknown tissue pH). Similar devices have been sold commercially and used in clinical practice, although not intended for perfusion monitoring, for example a combined pulse oximeter and $PtcCO_2$ sensor tested by Gisiger et al [90]. Furthermore, given the risk of permanent tissue damage from heated trans-cutaneous oxygen sensors, it would be desirable to use an unheated oxygen sensor for improved patient safety.

As described previously, trans-cutaneous oxygen sensors are usually heated to between 40°C and 45°C, to give a closer approximation to arterial oxygen tension, and enhance skin permeability. One factor that has not been described so far was explored by Stucker et al [91]; metabolic uptake of oxygen is significant in the dermis. In the presence of an oxygen sensor this implies that there will be oxygen consumption between the capillary bed and the sensor surface, as there will be no downward flux of O_2. Thus it can be seen that there are two potential obstacles to use of unheated $PtcO_2$ sensors; skin permeability needs to be sufficient, and the oxygen uptake in the upper skin layers must not be too significant as to lower partial pressure excessively (this is a less significant problem in heated trans-cutaneous O_2 sensing due to the increased perfusion and permeability resulting in increased oxygen availability). Eickhoff et al [92] found $PtcO_2$ correlation with blood pressure changed from negative to positive as the sensor temperature was elevated above body temperature. Peabody et al recorded $PtcO_2$ in neonates [93], and found trans-cutaneous O_2 dropped significantly below the arterial tension, but only in more severe cases of hypoxia/low perfusion. Durand et al [94] found strong correlation between trans-cutaneous O_2 measured on neonates with an unheated sensor and arterial PaO_2, but with trans-cutaneous O_2 partial pressure being significantly lower.

Taken together, it would seem that numerous studies confirm the conclusions that:

1. unheated trans-cutaneous partial pressure of Oxygen can measure tissue partial pressure, particularly in neonates,

2. the difference between tissue O_2 partial pressure and arterial pressure serves as a useful indicator of low perfusion.

However, it is unclear if $PtcO_2$ is sensitive to moderately compromised perfusion. An

"arterial to tissue dropout" metric also has no direct relationship to absolute perfusion, and may be challenging to use in a clinical setting due to unfamiliarity of clinical staff to the meaning of a newly introduced perfusion scale.

4.4 Conclusion

Trans-cutaneous blood gas monitoring is a diverse topic, offering several potential routes towards a non-invasive perfusion monitoring solution; both active (H_2), and passive (O_2 and CO_2).

Hydrogen based perfusion clearance techniques, involving hydrogen inhalation or generation by electrolysis within the local tissue, have been used in numerous studies, and are seen as one of the "gold standard" perfusion quantification techniques by many researchers. Unfortunately these involve invasive electrode implantation, causing considerable tissue trauma, and are thus unsuitable for routine clinical use in a continuous monitor.

Work by Varghese et al [62] shows that trans-cutaneous H_2 partial pressure measurement is feasible in adults, albeit with a heated sensor. Numerous papers indicate that the increased skin permeability and decreased epidermal thickness in neonates increases gaseous diffusion, and there would also appear to be potential to increase skin permeability through use of improved contact gels, suggesting that unheated $PtcH_2$ may be feasible in neonates, despite the fact the Varghese found it to be unreliable in adult patients. Varghese used a respirator and air supply with 2kPa partial pressure of added H_2, a very impractical system for use in continuous monitoring, even if the patient was already under supported ventilation.

It may be feasible to introduce H_2 trans-cutaneously using electrolytic H_2 generation within an adhesive patch sensor. The advantages of such a sensor over thermal techniques are debatable; problems of baseline temperature drift would be avoided, but the sensitivity would be affected by skin permeability, and air would have to be excluded from the sensor/skin interface much more effectively than in an equivalent thermal sensor. The design and manufacture of such a sensor would also appear likely to present difficulties; no existing published examples of a thin film hydrogen sensor with electrolytic H_2 generation could be found, and it would probably not be feasible to re-use commercial off-the-shelf trans-cutaneous oxygen sensors as previous researchers have if a means of H_2 generation was required within the sensor.

For these reasons fully trans-cutaneous H_2 clearance techniques seem somewhat unattractive as a route to achieving non-invasive perfusion monitoring, but more careful analysis

and modelling, combined with lab tests on volunteers may show this method to be more effective than thermal perfusion monitoring. It is notable that no studies could be found on trans-cutaneous H_2 diffusion into the skin. This is a factor that would need to be understood in more detail before a trans-cutaneous H_2 perfusion sensor could be ruled in or out.

In contrast to H_2 sensing, trans-cutaneous CO_2 and O_2 sensors are commonly used in clinical practice. There are a large selection of commercial off-the-shelf (COTS) instruments which could easily be employed as perfusion monitors if a suitable technique (i.e. either data processing or patient application related) could be devised.

A significant number of authors have found correlation between increased trans-cutaneous CO_2 partial pressure and low perfusion states, and trans-cutaneous CO_2 measurement has been found to be effective with an unheated sensor, an important feature as use of a heated sensor leads to risk of skin burns. $PtcCO_2$ is quite commonly used in existing clinical care, so use of peripheral and central unheated sensors for example may provide useful perfusion information without the introduction of new equipment or significant staff retraining. Existing research would appear to indicate significant variability in CO_2 partial pressure, one major contributory factor to which appears to be respiratory fluctuations. To counter this variability it would probably prove necessary to use either dual sensors (core and periphery) or perhaps an end tidal CO_2 monitor, although this would necessitate a breathing mask. Unlike H_2 clearance, CO_2 monitoring could not be used to quantify absolute perfusion, instead it would serve as a warning monitor to detect signs of falling tissue pH. The clinical effectiveness of $PtcCO_2$ as a low perfusion indicator is also relatively unproven compared to other techniques such as capillary refill timing and H_2 clearance.

Trans-cutaneous O_2 partial pressure has also been correlated to perfusion. Similar problems would seem to exist to those with CO_2; either a dual sensor configuration or some means to measure arterial oxygen tension would likely be required to obtain useful perfusion information. Pulse oximetry may be useful as a measure of arterial oxygen tension, but as arterial Hb is typically around 98% saturated, the sensitivity to partial pressure of O_2 will be low. It should be noted that it is precisely due to this low sensitivity of oximetry at high O_2 partial pressure that trans-cutaneous sensing is typically used.

It is also notable that use of dual site $PtcO_2$ monitoring has been demonstrated as a perfusion indicator, this is not the case with trans-cutaneous CO_2, but this does not appear to tell us anything about the feasibility of either scheme, as there are so few articles. The feasibility of unheated $PtcCO_2$ has been demonstrated by multiple authors, but the same cannot be said for $PtcO_2$. However this may be due to there being less

interest in unheated O_2 sensing, as it is widely believed that perfusion must be artificially elevated using heating for the tissue to be "arterialized" and the trans-cutaneous O_2 partial pressure to approximate that of the arteries. For perfusion monitoring we do not want this to be the case; we require tissue O_2 partial pressure, so the heating may well be unnecessary, but this has not been proven. The same points as with CO_2 apply to O_2 with respect to use only as a perfusion warning indicator.

In summary; trans-cutaneous monitoring of either O_2 or of CO_2 offers the potential for perfusion sensing, but this could likely only ever serve as a warning level metric, probably requiring either a second reference sensor or perhaps a respiratory sensor. There have been relatively few studies demonstrating the potential for such techniques to detect low perfusion, so more unequivocal demonstration of clinical effectiveness is required.

Chapter 5

Use of medical ultrasound for perfusion sensing

5.1 Perfusion sensing via ultrasonic imaging

5.1.1 Current clinical techniques in medical ultrasound

Medical ultrasound is a widely used technology with a long history. The first clinical usage being studies of the brain by Karl Dussik in 1942 (Paul G. Newman, "The history of ultrasound", [95]). In the 1950s, electronically scanned two dimensional medical ultrasonic imaging was developed by Douglass Howry and Joseph Holmes at the university of Colorado, and since then the field has expanded dramatically, with more recent developments including electronically scanned three dimensional ultrasound and Doppler ultrasonic scanning, often termed "colour Doppler" due to the method of visualisation.

The following nomenclature is typically used to distinguish the large variety of ultrasonic techniques;

- A-mode: the simplest and oldest technique, here a single directional transducer is used, and the reflected signal intensity versus time is plotted.
- B-mode: a linear array of transducers is used together with signal processing to image two dimensional "slices" through the body, with an image showing reflected signal intensity being displayed or printed.
- C-mode: constant depth mode using electronic time gating of the signal, often combined with mechanical scanning of the transducer and signal processing to produce an image.

- M-mode: any two dimensional technique where the image acquisition time is sufficiently short that video rate imagery can be produced. This enables diagnostic studies of organ motion, for example of the heart.

- Pulse inversion mode: a technique using two pulses of opposing phase, together with signal processing for detection of a non-linear reflection response. This non-linear compressibility phenomenon will typically arise if gas bubbles are present within the tissue or organ.

- Harmonic mode: in this mode non-linear tissue behaviour is employed for second harmonic generation. This technique is the subject of considerable recent research [96]. It is hoped that artefacts can be reduced by making use of the non-linear relationship between signal intensity and harmonic generation.

- Doppler modes: "colour" Doppler is typically used in modern clinical practice. This overlays two dimensional Doppler information onto a B-mode echo intensity display in the form of a colour scale. The intensity is displayed as brightness. On modern scanners, software permits blood vessels to be highlighted, and the total volumetric flow through the vessel to be calculated using numerical integration. Other information can also be displayed, such as a plot of volumetric flow rate versus time.

  Power Doppler employs a subtly different processing technique, whereby the total frequency shifted echo power is calculated, as opposed to the size of the frequency shift [97]. This technique has been used clinically to assess perfusion [98].

5.1.2 Prior studies of ultrasonic perfusion monitoring

Existing colour Doppler devices are suitable for spot checks of whole limb perfusion. A clinician can use the built in vessel flow integration software to assess the total blood flow into a limb through the principal arteries. However, such a technique would not be suitable for continuous monitoring, as the equipment is expensive and cumbersome. Additionally, only whole limb perfusion can be monitored, as only the larger principal arteries and veins have a sufficient diameter to be separated from surrounding tissue in the image. Use of colour Doppler for measurement of micro-vascular perfusion is challenging, as the velocities are small (of order 1mm/s or less). This means that the Doppler shift is usually below the sensitivity limit of the equipment. Hirsch found off-the-shelf colour Doppler was only suitable for distinguishing high perfusion states above 4ml/100g/minute [99]. This is around the typical volumetric perfusion found in

peripheral skin and muscle tissue, so colour Doppler would appear to be too insensitive to reliably detect reduced peripheral perfusion.

Use of the so called "power" Doppler signal processing algorithms as opposed to colour has been found to give a less artefact prone measure of perfusion [100]. However this technique still requires off-the-shelf two dimensional scanners, resulting in a solution that is cumbersome and expensive.

Injection of boluses of ultrasonically traceable contrast agents (for example micro-bubbles) into the arterial supply has been used in numerous studies, for example by Quaia et al [101]. However, this is invasive and clearly not suitable for continuous monitoring. For Doppler ultrasound to be practical in continuous monitoring, a cheaper, simpler, and less cumbersome technique than those currently in clinical use is required. Two dimensional imaging is not a requirement, especially if a technique can be devised that measures Doppler shifts arising from capillary bed flow.

Continuous wave Doppler has been proposed to meet these requirements, for example by Dymling et al [102]. However, Eriksson et al [103] found large variations in the measured perfusion, apparently due to the formation of speckle pattern. Jansson et al [104] discuss use of continuous wave Doppler with multiple closely spaced carriers as a method to mitigate the effect of speckle pattern. In practice, the extent to which this technique can be used is limited by the transducer bandwidth. Dylming et al successfully obtained perfusion measurements from a finger using a continuous wave device and a power Doppler processing technique [102]. The sensitivity to extremely low flow velocities ($<$1mm/s) was limited due to motion artefacts.

As Doppler shift is proportional to the vector dot product of the scatterer velocity and the scattering vector, a single emitter ultrasound device is limited to measuring velocity along a single axis. Numerous solutions to this problem have been proposed or are in use; integrated flow with colour Doppler will be unaffected by the angle between the scan plane and vessel of interest, as a vessel with a smaller angle will have a larger cross sectional area, and the integrated flow will always be the true flow. For imaging of the micro-vascular flow, where individual vessels are too small to be individually resolvable, assumptions regarding geometry can be made, usually an isotropic vessel structure is assumed (for example by Dymling). However, this assumption would appear questionable, as blood vessels can be significantly orientated in muscle tissue [105].

More sophisticated processing techniques using tomographic data have also been proposed, for example by Johnson et al [106]. These techniques make use of the fact that there are no isolated sources or sinks of blood within the vascular system, and therefore the vector divergence is zero, and it follows that the surface integral over any closed

volume will also be zero. Unfortunately, tomographic ultrasound must be employed for such techniques to become possible, and this inevitably results in an increased cost and size for the hardware applied to the patient.

5.2 Ultrasonically modulated optical tomography (USMOT)

5.2.1 An introduction to USMOT

USMOT is a considerably newer and less developed field than conventional ultrasound. Here ultrasound and optical techniques are combined, with the ultrasound being used to modulate an optical signal transmitted through tissue. Typically a long coherence length laser is used as a light source, with the ultrasound inducing phase modulation along ray paths scattered from tissue intersected by the ultrasonic beam. This modulation results in sidebands in the optical frequency domain, at frequencies of $\omega_l \pm \omega_u$, where ω_l is the laser frequency and ω_u the ultrasound frequency. At the photodetector face, the detected optical signal is proportional to the square of the optical electric field strength, and so the optical frequency sidebands produce intensity fluctuations at the ultrasound frequency ω_u.

The USMOT concept was first developed by Marks et al [107]. Wang et al developed the technique using time gated ultrasound pulses, publishing tissue phantom images obtained using single square-law detector [108]. USMOT in scattering media was analysed using the diffusion approximation by Leutz [109]. Both time and frequency domain USMOT has been employed, with early studies using pulsed ultrasound combined with time gating to obtain depth resolution, followed in 1998 by Wang and Ku [110] introducing frequency-swept ultrasound to obtain depth resolution along the ultrasonic beam axis.

5.2.2 Potential uses in perfusion measurement

A review of published USMOT research does not reveal any studies involving application of the technique to flow detection. Typically USMOT has been used to quantify tissue optical absorption μ_a at depth with relatively high spatial resolution, using the ultrasonically modulated signal intensity as an analogue of absorption.

Any optical paths through the tissue that also intersect the ultrasonic beam will also have a flow generated Doppler shift, caused by scattering from moving particles (i.e. erythrocytes). Considering this situation in optical frequency space, it can be seen that the ultrasonically generated sidebands will also have a flow induced Doppler broadening.

This Doppler shift will only have arisen from moving scatter centres intersected by the optical paths that pass also through the ultrasonic beam. Thus it can be seen that USMOT may offer a route to spatially resolved diffuse correlation spectroscopy (DCS), but the applicability of such a technique to perfusion assessment is unclear, and published studies in this area of application would appear to be lacking.

5.3 Photoacoustic effects

5.3.1 Photoacoustic effect in tissue

The photoacoustic effect in tissue arises due to thermal expansion caused by a small but very rapid temperature rise induced by optical absorption. Tissue is primarily comprised of water, and a consideration of the relatively high volumetric coefficient of thermal expansion ($2.07\times10^{-4}K^{-1}$) compared to the compressibility ($4.58\times10^{-10}Pa^{-1}$) suggests there will be a relatively strong thermally induced photoacoustic effect in water. For example, an instantaneous temperature change of 2°K will give rise to an instantaneous pressure rise of approximately 1MPa, sufficient to cause an intense acoustic amplitude, and even cavitation. Thus it would appear that the photoacoutic effect in tissue is sufficiently strong to cause damaging ultrasound levels at significantly lower optical intensities than those that would cause thermal injury, and so there is no potential risk of damaging tissue heating.

5.3.2 Use of photoacoustic effects for perfusion measurement

If a periodically pulsed laser is used to induce emission of ultrasound from absorbers via the thermally induced photoacoustic effect, then any flowing absorber will emit ultrasonically at a Doppler shifted frequency. Thus an ultrasonic transducer placed in contact with the skin surface will be capable of detecting moving absorbers (i.e. blood flow). Unfortunately, although flow velocity information can clearly be obtained, total flow volume (i.e. perfusion) is not clearly related to the measured Doppler shifted ultrasound. Attempting to derive absolute flow from the absolute amplitude of the Doppler shifted ultrasound is complicated by the unknown optical coupling coefficient into the skin, unknown amount of stationary photoacoustic emission within the tissue, and the unknown acoustic coupling coefficient from the skin surface into the transducer.

The use of this technique to obtain flow information from moving scattering particles was first proposed by Fang et al in 2007 [111]. In their published results, the authors demonstrated the detection of low flow velocities of <0.1mm/s. In a separate paper,

similar results were obtained in a scattering phantom [112]. A three dimensional velocity scanning system is described by Yao et al [113], but this technique does not appear to be suitable for use in scattering media. Brüker et al [114] demonstrated velocity measurements from a rotating phantom consisting of micro-dots. Zhang et al [115] demonstrated use of a multi wavelength photoacoustic technique to determine the oxygen saturation of capillary vessels in vivo.

Use of the photoacoustic effect to accurately quantify blood perfusion in tissue has not been demonstrated. However, comparatively few data have been published in this area, and with low cost diode lasers now available that can generate peak powers in the hundreds of Watt range, the photoacoustic may offer a relatively low cost and simple to implement route to perfusion assessment.

5.4 Conclusion

Used appropriately, ultrasound is a non-ionising, non invasive and safe diagnostic tool, that does not result in tissue damage. It is already in wide clinical use, and off-the-shelf devices are capable of providing useful spot checks of whole limb perfusion. Unfortunately, off-the-shelf ultrasound devices are not capable of reliable measurement of the low flow velocities found in the capillary bed, and are too cumbersome and expensive to be a practical solution for continuous monitoring. There has been extensive research into continuous wave Doppler ultrasound for measurement of low flow velocities in the perfused capillary bed, but no commercial solution has ever entered the market. Artefact problems from tissue motion are significant, due to patient movements and respiration causing Doppler shifts of the same magnitude as the perfusion shifts. Existing commercial devices (e.g. 2D colour Doppler flowmeters) use Doppler threshold filters to remove artefacts, but these techniques would not be applicable for slow capillary bed flow.

An additional complication with all ultrasonic techniques is the requirement for effective coupling to the tissue. Coupling is significantly more challenging with ultrasound than with optical and thermal techniques due to the very significant mismatch in acoustic impedance between air and tissue. The majority of production devices are hand-held, and used with a coupling gel to provide impedance matching between the skin surface and the transducers. For a continuous monitoring device to be successful, problems of coupling would need to be mitigated, which may complicate device fixation in a medical environment.

Hybrid optical/ultrasonic techniques (USMOT and photoacoustic) are a relativly new and emerging field in bio-medical imaging. The applicability of USMOT for flow analysis is in doubt, but photoacoustic Doppler has been demonstrated as a method of flow sensing (but not perfusion) in phantom studies.

Chapter 6

Optical perfusion sensing

6.1 Introduction

Optical techniques form a large and well developed area in biomedical monitoring. In this chapter, five areas of interest are explored in detail. There is some crossover between optical techniques and methods discussed in other chapters, such as optically induced heating for use in thermal monitoring, and the photoacoustic effect discussed in the previous chapter. This chapter concentrates on purely optical techniques, and techniques involving external pressure application combined with spectroscopy.

6.2 Laser Doppler

6.2.1 Reflection mode laser Doppler/Quasi-elastic light scattering (QELS)

Reflection mode laser Doppler records variations in the speckle pattern formed from coherent laser illumination of the skin surface. This technique is in widespread use, and numerous manufacturers offer off-the-shelf hardware. The term "quasi-elastic light scattering" is often used to describe this technique, particularly in the case of diffuse illumination, where photons are typically single scattered back out of the tissue before reaching the detector. Both single point adhesive devices capable of outputting a continuous measurement, and scanning devices, capable of producing a two dimensional image of the skin surface are available.

Measurement of volumetric perfusion is not attempted in these systems, instead a simple "perfusion metric" is output. This is sometimes defined as the AC/DC ratio of the speckle fluctuations, where the dividing line between AC and DC is some arbitrary frequency point in the range of tens of Hz [116] (20Hz is common). Defining the speckle

modulation power at angular frequency ω as $P(\omega)$, and PI as the perfusion index (a metric that may be output from the monitor), the following equation is applied to the raw intensity data by the laser Doppler device:

$$PI = \frac{\int_{\omega_{ac}}^{\omega_n} P(\omega)d\omega}{\int_0^{\omega_{ac}} P(\omega)d\omega}. \tag{6.1}$$

In equation 6.1, ω_{ac} is the frequency separating the "AC" and "DC" signal components, and ω_n is either the Nyquist frequency of the photodetector system or some upper limit in the range of tens of kHz. Holloway states that this upper frequency ω_n is usually chosen to be in the order of 20kHz to 1MHz [116], but device manuals available from laser Doppler device manufacturers reveals that in practice, 20kHz is a more common choice. Alternatively in an analogue system, a low pass filter may be added to exclude high frequency noise, followed by further filtering to isolate the "AC" and "DC". In digital systems, a low pass stage is often added to the front end to exclude shot and thermal noise at frequencies above the typical speckle fluctuation range from entering the ADC. The first moment of the speckle fluctuation power spectral density is also commonly used as an approximate measure of perfusion. This is defined as (normalised with respect to intensity)

$$PU = \frac{\int_{\omega_{ac}}^{\omega_n} P(\omega)\omega d\omega}{\int_0^{\omega_{ac}} P(\omega)d\omega}, \tag{6.2}$$

and is the metric used by the Moor Instruments MoorVMS single point laser Doppler unit, amongst other devices. The PU coefficient is found to be approximately proportional to perfusion [117, 118], but is highly dependent on geometry of the vasculature and micro-vasculature [119, 120], and so cannot be standardised between patients or measurement sites. However, it has been found to give a reproducible measurement of perfusion trends in many studies, both between patients and between measurement sites [121, 122].

More sophisticated analysis of reflection mode laser Doppler data is difficult due to the complex relationships between the blood flow and the laser speckle. These problems can be seen from a consideration of typical two dimensional scanning laser Doppler images - principal blood vessels can be seen as very significant local increases in the PI index, which is highly heterogeneous.

6.2.2 Diffusing wave spectroscopy and Diffuse correlation spectroscopy (DWS/DCS)

DWS attempts to resolve some of the heterogeneity issues of reflection mode laser Doppler by increasing the path length and fitting a theoretical model to the auto-correlation function of the laser speckle pattern data. Bonner et al [123] were the first authors to analyse coherent light propagation in a mixture of moving scatter centres. Mariet and Wolf [124] further developed this theoretical basis, applying the diffusion approximation to the case of a highly scattering medium.

The system of equations developed by Mariet and Wolf included a geometry dependent term, a function of source and emitter placement relative to the tissue, which was itself modelled as a heavily scattering semi infinite half space. This means it is extremely difficult to apply the Mariet and Wolf model to real world situation. However, in the case of diffuse coherent illumination and an infinite scattering slab, Brown [125] demonstrated that the following equation describes the intensity autocorrelation function $g_2(\tau)$ in the centre of the illuminated area.

$$g_2(\tau) - 1 = exp\left(-2\gamma\sqrt{\langle\Delta r(\tau)^2\rangle k_0^2}\right), \tag{6.3}$$

where the coefficient γ is found to have a value of approximately 2, and k_0 is defined as:

$$k_0 = \frac{2\pi n}{\lambda} \tag{6.4}$$

with n equal to the tissue mean refractive index.

Boas et al [126], and Mason et al [127] developed a system of equations to describe inter-speckle autocorrelation in DWS. In 2001, Cheung et al [128] developed these concepts into an in-vivo DWS device using temporal correlation between two physically separate detection fibres. The use of dual detectors gave increased immunity to artefacts caused by patient motion, and reduced sensitivity to blood vessels in the immediate vicinity of the detection fibres, resulting an a more accurate measure of average volumetric flow. This technique is often termed "Diffuse Correlation Spectroscopy" (DCS), although the nomenclature used to distinguish between laser Doppler techniques can be inconsistent in the literature.

As the optical scattering characteristics of erythrocytes are well characterised and relatively consistent between individuals, it would appear that DCS allows measurement of $\langle\Delta r(\tau)^2\rangle$ to be interpreted as a volumetric perfusion of red blood cells. Blood flow would apparently cause $\langle\Delta r(\tau)^2\rangle$ to take a different form from that of Brownian thermal

motion. This can be seen by considering the expectation value of particle displacement $E(\Delta r(\tau))$, which will be proportional to time τ for flow, but to $\sqrt{\tau}$ in Brownian motion (i.e. a random walk).

However, in in-vivo studies, it has been found that only a Brownian form gives a good fit to the autocorrelation function. Zhou [129] found $E(\Delta r(\tau))$ took a fully Brownian form in both muscle and brain tissue. Yu et al [130] obtained DCS data from muscle in-vivo, with simultaneous MRI Arterial Spin Labelling (ASL), which is clinically accepted as a gold standard perfusion metric. A Brownian diffusion model was fitted to the DCS data, and the best fit diffusion coefficient plotted against ASL perfusion. A strong correlation was found between DCS and ASL, but relatively significant variability in the DCS derived diffusion coefficient at low perfusion values (<30ml/100g/min) implied that DCS was not an accurate measure of low perfusion states.

The reason that perfusion results in scatter centre motion closely resembling a random walk is unclear, but it is hypothesised that high shear rates and turbulent flow are responsible [131]. A single red blood cell contains multiple scatter centres, so tumbling motion in a turbulent flow would likely resemble a random walk of scatter centres rather than linear motion.

6.2.3 Doppler coherence techniques

In order for an optical frequency Doppler shift to be observable, the combining scattered light paths must have a differential pathlength of less than the laser coherence length, or they will be incoherent when they recombine, and no interference will occur. Using a reference beam to control this coherent interference in techniques similar to optical coherence tomography (OCT), information on the relationship between Doppler shift and path length can be derived. The application of coherence tomographic techniques to optical Doppler was first explored by Chen et al [132] in 1997. Bizheva et al [133] (1998), and Wax et al [134] (2001) explored a similar technique, using longer path lengths in phantoms to investigate the transition to multiple scattering, and there have since been further replications.

The technique developed by Chen et al allowed three dimensional images to be obtained from rat skin by use of a two axis scanning confocal microscopy platform. Image slices were taken at a depth of $200\mu m$ below the skin surface. This is clearly a complex technique to implement, and not well suited to routine clinical monitoring. The scan depth is limited to of order one transport mean free path (i.e. $1/\mu_s'$), meaning only a very local measurement of perfusion close to the skin surface can be obtained. Liu et al [135] obtained three dimensional images of vascular flow at depths of up to 1.2mm

during in-vivo trials in human skin.

It may prove feasible to replace the two axis scanning system with a CMOS imager and beam splitter based illumination system. This would result in a comparatively simpler design, but to replicate the Chen method exactly would require an extremely high frame rate imager.

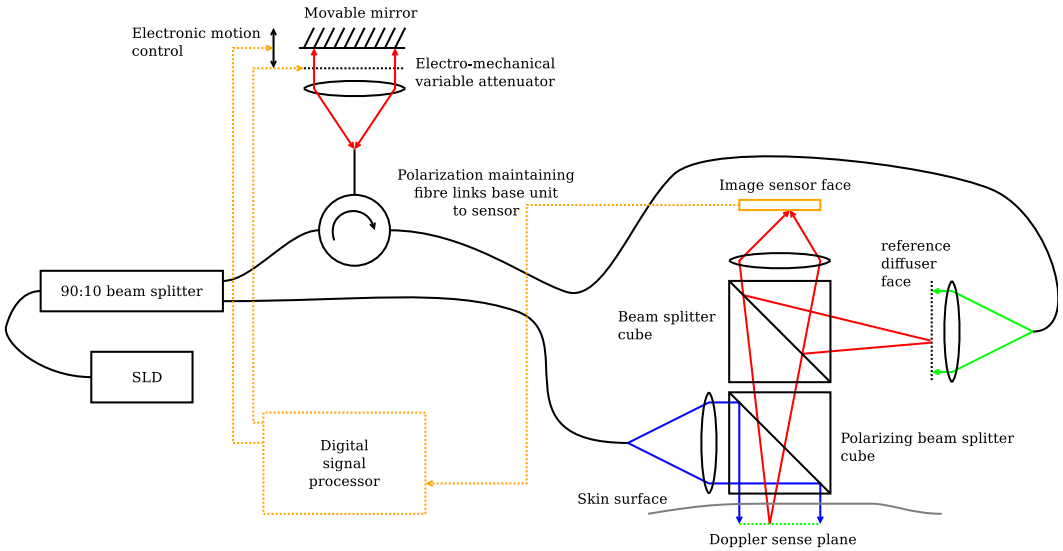


Figure 6.1: Author's proposal for a Doppler OCT device

Doppler OCT suffers from the same problems of turbulence induced red blood cell tumbling as DWS and DCS techniques, but through use of optical coherence, Chen et al were able to obtain a full frequency spectra of a small voxel of tissue. By use of a model fitting algorithm, it was then possible to calculate the flow velocity in the vertical z axis. A CMOS imager would require a frame rate of around 10kHz or more to replicate this technique, which would be prohibitively expensive (such sensors are sold commercially for prices in the range of several thousands US Dollars) with currently available sensors, and require an extremely high bandwidth data path from the sensor to the DSP. An alternative route was taken by Choi et al [136]. By amplitude modulation ("strobing") of the light source it was possible to obtain the full Doppler spectrum from a limited frame rate camera. This allows use of extremely low cost CMOS imager sensors, which mostly have a frame rate of 60Hz or lower. Due to the widespread adoption of CMOS cameras in consumer electronics, these devices are less than an order of magnitude more expensive than equivalent area photodiodes, and the cost per unit seems likely to decrease further in future.

Zhang et al [137] used a swept frequency laser source in place of a low coherence superlumescent laser diode (SLD). It may prove feasible to combine swept frequency with the apparatus illustrated in figure 6.1, which would result in an improved acquisition time if multiple z slices of Doppler information were desired.

6.3 Near-Infrared spectroscopy (NIRS)

NIRS in tissue is a large field, with multiple techniques in use or development, and several manufacturers marketing off-the-shelf hardware for clinical use. Unlike Doppler techniques, spectroscopy cannot give any direct measurement of flow, but it may still provide useful information. For example the Fick principle [138] allows the difference between venous and arterial oxygen saturation to be related to volumetric haematocrit perfusion if the metabolic rate of the tissue is known. Lima et al [139] found NIRS to be an effective indicator of low peripheral perfusion.

Tissue is highly inhomogeneous; the difference in optical absorption coefficient μ_a can change by orders of magnitude over both short (μm) and long (mm) scales in the tissue. In the near infrared from approximately 700nm to 950nm, it is found that the absorption μ_a is relatively low (leading to the term "near infrared window" to be used to describe this range of wavelengths), meaning that inhomogeneities in tissue absorption tend to add in a linear fashion, and so it becomes feasible to employ homogeneous approximations based on the average absorption coefficient. Spectroscopy is also complicated by the strong optical scattering in tissue. However, within the near infrared window region, the transport corrected scattering coefficient μ_s' is found to be around an order of magnitude higher than μ_a, making transport models derived from the diffusion equation valid. Additionally in this region μ_s' is found to follow a power scaling law with wavelength that is relatively consistent between tissue types. It is for these reasons, combined with the large penetration depth that comes as a result of the low absorption, that spectroscopy of bio-materials is usually carried out within the near infrared window.

In order for spectroscopy to be quantitative, the optical path length through the tissue must be known or estimated. It is the method of obtaining the distribution of path lengths (also commonly referred to as the Temporal Point Spread Function or TPSF) that distinguishes the principal methods of NIRS in tissue.

6.3.1 Continuous wave

In terms of required hardware, continuous wave is the simplest NIRS method. Using several incoherent light sources (such a Light Emitting Diodes or LEDs), which are either unmodulated or slowly modulated with respect to the TPSF (in order that sources may be separated using signal processing methods), continuous wave techniques use multiple optical wavelengths, or physical source to detector separations as a method of path length estimation. Scholkmann et al give an overview of current developments in the field in [140].

If the diffusion approximation is used to model photon transport, and at least two combinations of source to detector separations and optical wavelengths are used, then the ratio of the tissue absorption coefficient to the transport corrected scattering coefficient $\frac{\mu_a}{\mu_s'}$ can be derived without any assumptions concerning optical coupling efficiency and skin melanin absorption (a principle termed "self calibrating" by Heuber [141]). If the tissue saturation i.e. the overall fraction of tissue haemoglobin in the oxygenated state is desired, then a power law approximation can be applied to estimate μ_s', and tissue saturation can be found without knowledge of the absolute value of μ_s'. This technique was found to produce saturations with a good match to those obtained via blood sampling by Matcher [142].

Tissue saturation spectrometers are available off-the-shelf, for example from Artinis Medical Systems B.V., Zetten, Netherlands [143]. Trends in tissue saturation have been correlated with changes in perfusion [144], but absolute perfusion cannot be obtained via tissue saturation data, and a spot check of tissue saturation is of only limited use, but can serve as an indicator of seriously compromised blood supply [145].

Simpler techniques for measurement of changes in HbO_2 and Hb using multiple wavelengths and a single separation have also been published. Matcher et al [146] compared the performance of several different differential path length algorithms. These assume that if the measured change in tissue absorption is small, the absorption change will be linear with absorber concentration change, and hence by taking the first derivative of absorption with respect to absorber concentration, the problem can be linearised. Effects of scattering and baseline absorption can be approximated as a so called "differential path length factor" or DPF, which will be equal to the ratio of the mean free path between source and detector to the physical separation. Remarkably consistent DPF values have been obtained from the forearms of adult volunteers in numerous studies, for example by Duncan et al[147], Kohl et al [148], and Essenpreis et al [149]. The DPF is found to typically equal around 4.2 at 800nm, with a moderate and approximately linear decrease as wavelength increases. Kohl also found that the wavelength dependence of the pulsatile signal amplitude was an effective method of determining the change in DPF with wavelength, but not the mean DPF. Clearly these DPF techniques can only be applied in the case of changing haemoglobin, as they cannot give useful information on the baseline state as source intensity cannot be compensated for, but they will be revisited later in this chapter in the context of processing data obtained via techniques such as capillary refill.

6.3.2 Modulated

Amplitude modulation of the light sources at frequencies in the region of 110MHz is a technique used in many devices. The convolution of the input sinusoidally modulated intensity function with the tissue TPSF leads to both a phase lag and a decrease in modulation depth. It is not possible to obtain the full form of the TPSF from this method, but with knowledge of the source to detector separation, the phase lag and modulation depth can be fitted to a diffusion model giving tissue μ_a and μ_s'.

The use of relatively low frequencies and narrow bandwidth with respect to time resolved techniques makes modulated light attractive. It is in use in numerous off-the-shelf devices, for example Hutchinson Technology market a range of spectrometers [150] using modulation at 110MHz.

The assumptions made in this technique are only applicable to homogeneous media. Housley [151] found that the diffusion model broke down in the cases of inhomogeneous or layered tissue.

6.3.3 Time resolved

Time resolved techniques have typically used a pulsed light source together with a streak camera or photomultiplier tube. Avalanche photodiodes, and single photon avalanche diodes (SPADs, photodiodes operated in Geiger mode above their breakdown voltage), have also been employed. In typical tissue spectroscopy applications, with a source to detector separation measured in centimetres, a temporal resolution of tens of picoseconds is required for a useful TPSF resolution to be obtained. This implies a bandwidth in the tens of GHz, a challenging engineering requirement.

Large photomultiplers and high voltages mean that in most systems fibre bundles are used between a base unit and a patient mounted sensor, for example the "MONSTIR" device constructed at the UCL BORL laboratory, and described by Cooper et al [152], which uses a large trolley mounted base unit containing a pulsed super-continuum laser, photomultiplier tube array and optical switching and attenuator matrix. A large number of fibre bundles run between the base unit and patient mounted sensor assembly, and the switch matrix permits all possible combinations of source and detector locations to be selected. Inversion algorithms allow for full three dimensional reconstruction of tissue μ_a and μ_s'.

Currently, fully time resolved systems are extremely expensive and cumbersome, and appear ill suited for use in continuous monitoring. Use of avalanche photodiodes in place of vacuum tube based photomultiplier tubes or streak cameras may allow use of patient

applied light detection, and so avoid the use of large fibre bundles, but as avalanche photodiode gain is much lower than that of photomultiple tubes (typical sensitivity is no better than 10^2 photons [153]), obtaining a high SNR measurement of the TPSF with avalanche photodiodes is challenging. Avalanche photodiodes also require high voltages (of order 60V) to be present within the patient applied sensor enclosure. This could present potential hazards and regulatory issues.

SPAD devices have a relatively long recovery time after triggering (longer than the TPSF in tissue), meaning they are typically employed as time correlated single photon counters, requiring many pulses to be emitted from the light source before an accurate TPSF (i.e. one with high SNR) can be produced. To maximise SNR, it is also necessary to reduce the dark count rate to low levels in SPADs, leading to only small area devices being practical for use in time correlated single photon counting. Such devices are poorly suited for use as diffuse light detectors in NIRS tissue studies, although SPAD based NIRS devices have been demonstrated in in-vivo studies using small source to detector separation (e.g. by Di Sieno et al [154]), and relatively large area (200μm diameter) SPADs with low dark count rates (1500/s) have been produced by Ghioni et al [155], although this device was cooled to -25°C. SPAD arrays, with many densely packed SPADs giving a larger effective collection area, have the potential to resolve these difficulties. Arrays with total collection areas as large as 1mm^2 have been fabricated successfully [156]. Recently (2014), such integrated SPAD array devices have started to appear on the market, some with light source drive and digitisation electronics integrated on die, for example, although ST microelectronics markets the VL6180X as a proximity sensor for applications in consumer electronics [157], the sensitivity of this device is sufficient for use in tissue NIRS studies.

It would appear that the very low cost SPAD array devices currently making their way into consumer electronics applications may offer the potential for cheap, practical, time resolved NIRS to become readily available in the near future.

6.3.4 Coherence and optical frequency resolved

The use of optical coherence techniques has already been discussed in the context of Doppler flow techniques, and is also commonly employed in three dimensional optical tomography of tissue via optical coherence tomography (OCT). These methods are all used over relatively few mean transport free scattering depths, in other words they are single scattering techniques. However, Tualle et al have extended the use of optical coherence to the full scattering limit [158], where both temporal [159] and frequency domain optical coherence were used to obtain the full form of the TPSF. Results were

obtained both with fibre coupled photodiodes for detection and with a free space CCD speckle contrast imaging system and variable coherence length laser illumination [160]. It should be noted that Tualle holds a patent on the use of fully scattering coherence techniques to obtain the TPSF in tissue [161].

Use of optical coherence in tissue presents practical difficulties, as in the case of the diffuse (i.e. subtending of order 2π steradians solid angle) light leaving the surface of a fully scattering material, the laser speckle grain size is very small, of the same order as the optical wavelength. Although it is possible to use optical assemblies to increase the speckle size, this comes at the expense of brightness, and it is thus apparent that there is a fundamental limit to the amplitude of the detectable speckle fluctuations.

In addition to the modulation induced speckle, there are also flow induced optical Doppler shifts in tissue, leading to strong speckle fluctuation on timescales of milliseconds or less. Tualle suggests an optical frequency ramp function with a period in the microsecond range, moving the differential path length interference out of the Doppler frequency band [160], but this means photodiode detectors are essential in order to have sufficient bandwidth. Use of a larger photodiode only increases the size of the speckle signal with the square root of the detector face area, meanwhile the overall detected signal increases linearly with area, so although the speckle signal amplitude can be increased with a larger detector, it is rapidly overtaken by power fluctuations in laser output power. Published studies have used acousto-optic modulators and balanced detection to mitigate this problem, but the speckle signal strength has still been low. Use of high speed optical frequency modulation to simulate a variable coherence source, as described by Tualle in [160] may be another route around this problem. Speckle contrast imaging with an off-the-shelf CMOS image sensor at variable simulated coherence lengths would allow the full form of the TPSF to be computed.

Commercial swept frequency laser systems are custom built high cost items, but there has been considerable interest in obtaining single frequency (i.e. single longitudinal mode) operation from mass produced edge emitting Fabry-Pérot laser diodes within the hobbyist laser holography community. Using precise temperature and current control, single optical frequency mode is possible with some commercial laser diodes, which are usually multi mode. Single longitudinal mode operation has been obtained from the Mitsubishi ML101J27, Hitachi HL6385DG, and Opnext HL63133DG ([162] and personal email correspondence). ECDL designs are also feasible with many off-the-shelf laser diodes [163]. With all off the shelf laser diodes, coupling to a single mode fibre is necessary for optical coherence techniques to be possible. This is probably the most challenging problem if such laser diodes are to be used.

6.4 Use of venous and arterial occlusion with NIRS

Pressure cuffs have been used for many years in volume plethysmograhy studies. Applied to limbs, these block the return of venous blood (with a low application pressure) and inflow of arterial blood as well (higher applied pressure). In its usual form, used together with a mechanical limb volume measurement technique (such as a water bath) this technique is somewhat limited in that it can only measure the total blood flow to a limb.

More recently, NIRS has been applied to pressure cuff studies as the means of measuring blood volume. This technique has a major advantage in that with venous occlusion the blood accumulates locally [164], permitting a direct measurement of local perfusion from the rate of total haemoglobin increase. The venous saturation can also be calculated by measurement of the ratio of oxy-haemoglobin to total haemoglobin rate, and assuming that blood accumulation takes place primarily within the venous compartment [165]. This would seem reasonable given that veins and venuoles are much compliant than arterioles.

An alternative means to measure venous saturation involves use of arterial occlusion to completely isolate the blood within the limb. The rate of desaturation can then be used to calculate the metabolic rate, and the venous saturation found from the perfusion by use of the Fick principle, a technique outlined by Beekvelt [164].

Venous occlusion has been widely employed in many published studies, and found to correlate well with perfusion measurements obtained via "gold standard" methods such an arterial spin labelled MRI. For example by Raynaud et al [166], and Woinarski et al [167]. This makes it perhaps the most appealing non-invasive perfusion monitoring technique so far discussed, having a significant research base of published studies, being completely non-invasive and potentially (in the form of continuous wave NIRS) involving low cost equipment. It is also notable that only differential measurements of Hb and HbO_2 are required for occlusion techniques to be effective, making the problem of inverting the measured intensity data to haemoglobin molarities significantly easier. Specifically, if the DPF algorithm is to be employed, a simple multi-wavelength single separation device can be used with no calibration required for light source intensity and coupling into the tissue etc.

6.5 Direct imaging of capillary flow (Capillaroscopy)

Direct imaging of the upper capillary loops is possible with simple microscopy techniques in areas of extremely thin epidermis, such as the nail bed and inner lip. With sufficient magnification, transit of individual erythrocytes through the microcirculation can be imaged directly [168]. However, for imaging in areas of more typical skin thickness, where the upper capillaries are obscured by a much more heavily scattering layer, simple microscopy techniques are insufficient to resolve capillary flow. Dobbe et al [169] developed automated software algorithms to calculate capillary flow velocities, functional capillary densities and transit times from video capillaroscopy. Theoretically this information would allow a volumetric perfusion value to be derived, but only for the upper layer of capillary loops.

It should be noted that Capillaroscopy is perhaps more suitable for neonatal use than for use on adults, as adult skin is considerably thicker and thus more heavily scattering.

6.5.1 Orthogonal polarization imaging (OPS)

This technique is discussed in [170]. By illuminating the skin surface with polarized light, it is possible to remove photons that are single scattered from the skin surface or upper layer of skin. Multiple scattered photons lose their polarization, and can pass through the orthogonal polariser placed over the detector to form an image of the deeper tissue structures. This is found to give improved image quality over conventional microscopy [171].

6.5.2 Dark field microscopy in skin

Dark field microscopy involves arrangement of microscopy illumination such that only light scattered from the sample is visible in the microscope image. The technique is typically applied to capillaroscopy in the form of "side-stream darkfield imaging" or SDF. The SDF technique uses a ring of light sources surrounding a circular central region of skin where the microscope is focussed. Usually the illumination is stroboscopic, and may use multiple light emitting diodes at different wavelengths to obtain blood oxygen saturation spectroscopically. Light entering into the tissue via the outer illuminated ring is scattered multiple times before exiting into the microscopy optics. The majority of the observed light having scattered to several mean transport free depths before exiting upwards towards the microscope. Thus it can be seen that SDF is not a truly dark field technique, as the background is actually heavily illuminated by scattering light. Rather,

the observed blood vessels are back illuminated by the scattered light, and thus appear as dark outlines on the image. Goedhart et al found stroboscopic SDF gave moderately improved results over OPS [172].

6.5.3 Optical coherence tomography (OCT) techniques

OCT has already been discussed in the context of Doppler tomography and NIRS, but OCT may also have the potential to track flow of individual cells during capillary transit as part of an improved performance capillaroscopy system. In-vivo OCT of the human skin was demonstrated by Welzel et al [173]. White et al [174] demonstrated use of high speed OCT for direct observation of capillary flow in the retina, and Barton et al [175] demonstrated flow measurement in phantoms containing intralipid. Enock et al [176] obtained flow information in mouse brain using a similar technique, and Doron et al [177] presented GPU assisted algorithms for characterising the vascular networks from OCT images of in-vivo human skin. Yousefi et al [178] developed a system dubbed "Ultra High Sensitivity Optical Microangiography" or UHS-OMAG, a hybrid Doppler OCT technique using Doppler in the depth or z axis, and correlation in the x and y axes, to give a three axis flow estimate in the in-vivo mouse ear.

In a thorough literature review, no published implementation of a pure correlation OCT flowmetry system in in-vivo human skin has been found. This may well be due to the attractiveness of Doppler-OCT compared to capillaroscopy. Doppler can operate with relatively poor spatial resolution; sufficient to resolve the individual vessels, whereas capillaroscopy flow techniques require resolution sufficient to resolve individual red blood cells. Doppler is however limited to resolving flow velocity along the depth (or z) axis. Capillaroscopy is capable of resolving two axis flow, but calculation of volumetric flow is difficult, as the image is only two dimensional. High spatial resolution OCT techniques may make accurate determination of volumetric flow possible, but coherent Doppler methods would appear to be a more promising route to take if this was desired.

6.6 Photo-plethysmography (PPG)

Plethysmography is the measurement of volume changes in the body or an organ. Photo-plethysmography develops this method into a more practical form through use of optical absorption to measure changes in blood volume. Typically the pulsatile signal from the arteries is filtered and used to estimate heart rate and oxygen saturation (via pulse oximetry). However, blood volume also fluctuates with the respiratory cycle, temperature regulation, and motion of the body, leading to a venous component to the PPG

signal that may also provide useful information.

6.6.1 PPG amplitude and perfusion

The pulsatile modulation depth has been suggested by many as a perfusion metric. It is currently measured by several commercial pulse oximetry systems, and is marketed as a measure of perfusion viability. Lima [14] discusses its use in critical care and healthy adult volunteers, including multiple studies where "Peripheral perfusion index" or PI was recorded using commercial off-the-shelf pulse oximeters from Philips Healthcare, and was found to be highly variable amongst a large group of healthy volunteers. However in the critical care group trends in PI were found to correlate with changes in core to peripheral temperature difference, suggesting that PI is an effective means to measure changing peripheral perfusion. Felice et al [179] found that PI trend was an effective means to assess illness severity in neonatal care. Zaramella et al [180] found correlation between calf muscle perfusion measured using NIRS combined with venous occlusion and PI measured on the foot in a group of 43 healthy neonates, but the PI was highly variably between individual subjects, with a correlation analysis between NIRS derived perfusion and PI yielding a Pearson r coefficient of $r = 0.32$ with $p = 0.03$ and thus suggesting that a spot PI assessment is relatively meaningless for determining if the limb was adequately perfused, despite the fact that the PI correlated significantly to perfusion.

Nevertheless, the authors commented that in their opinion PI was a more practical means of perfusion assessment than NIRS derived tissue oxygen saturation, both due to the practicality of the equipment (i.e. an off-the-shelf pulse oximetry system that would likely already be in use and familiar to medical staff) and the quality of the PI data.

6.6.2 Extracting venous saturation

As discussed previously, not all the observed PPG signal is arterial in origin; there are signal components arising both from the different vascular compartments, such as the capillary bed and the veins/venuoles, and from changes in the tissue structure due to for example motion induced compression.

Masimo Inc. have patented a novel signal processing technique for use with pulse oximeters. U.S. patent number 5,482,036, 1996 [181] describes a method whereby adaptive noise cancelling is used to isolate the arterial signal component from the pulse oximeter signals. The arterial blood oxygen saturation can then be found, as well as venous saturation in some cases. According to promotional material from Masimo, the so called

"Discrete Saturation Transform" (DST) algorithm can outperform conventional pulse oximeters in cases of motion artefact. There have been a few peer reviewed papers published on the performance of the DST approach, and controversy about the true accuracy of the data obtained from DST, but only one journal article [182] describing the details of the patent could be found.

The fact that DST can allow venous blood oxygen saturation to be determined in cases

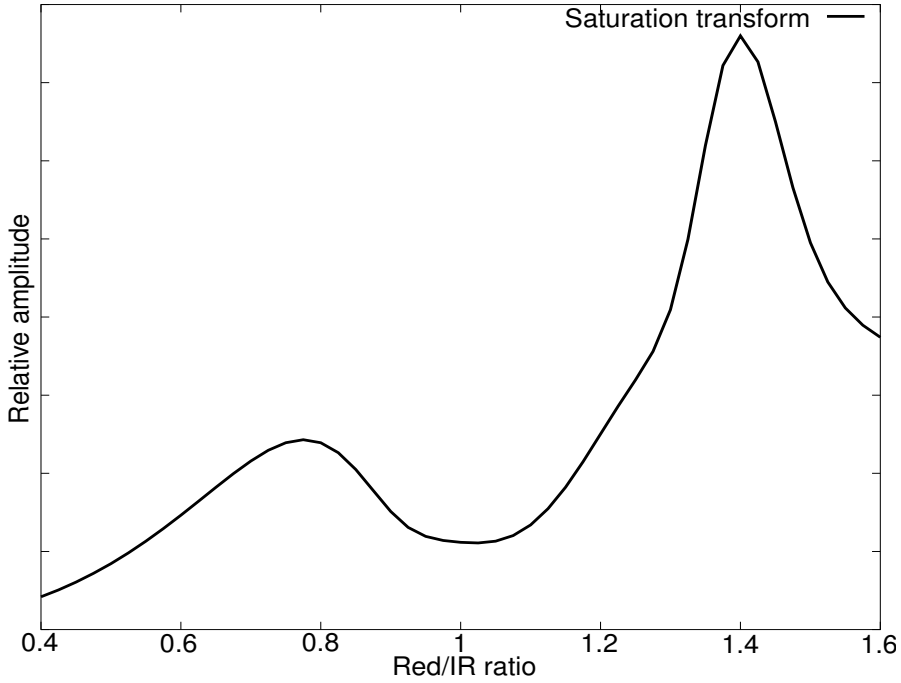


Figure 6.2: Masimo "Saturation transform" algorithm output

of motion artefact is interesting from the point of view of perfusion monitoring, as low perfusion may lead to excessively low venous saturation (from the Fick principle). The percentage of available blood oxygen extracted during blood passage through the capillary bed, or oxygen extraction index, is related to perfusion due to the oxygen demand of the tissue being roughly constant with time. There is potentially a problem with use of DST to extract venous saturation when no motion artefact is present, but an automated capillary refill device could induce a venous PPG signal, allowing a Masimo type algorithm to determine oxygen extraction index even when no motion artefact is present, a concept patented in Chan in [183]. There is also evidence to suggest that the respiratory cycle may induce bulk motion of venous blood [184].

Zaramella et al [180] found no significant correlation between PI measured on the foot of a group of neonatal patients and oxygen extraction index measured via NIRS and venous occlusion, but they did not check for a correlation between perfusion and oxygen extraction, and the poor correlation between PI and perfusion seen in their study could have lead to such a correlation being missed.

However, significant (Pearson $r = 0.32$, $p < 0.05$) correlation between the NIRS derived delivered oxygen and the PI was found. This is interesting as it supports the conclusion that perfusion is regulated to ensure sufficient supply of oxygen to the tissue. Therefore a means to detect increases in oxygen extraction index seems likely to offer a route to accurate detection of perfusion adequacy.

The Masimo algorithm was implemented in Matlab using least mean squares for the adaptive noise canceller (the noise cancelling adaptation technique is not described in the patent). Simulated Red and Infra-red data were generated using sinusoidal "pulsatile" and "motion artefact" signals of differing frequencies. The Red/IR absorption ratio for venous and arterial blood were set at 1.4 and 0.75, corresponding to oxygen saturations of approximately 70% and 95% saturation (for example see [185] for pulse oximetry calibration curves).

From Figure 6.2 it can be seen that the Matlab code produces two peaks in filter output root mean squared amplitude, at approximately the same red/infra-red ratios as used to create the input data.

6.6.3 Multi-wavelength PPG and tissue oxygenation

There is no obvious distinction between multi-wavelength photo-plethysmography and near infra-red spectroscopy techniques based on changes in optical absorption. In the following section, the term PPG will be used solely to describe use of the pulsatile signal component in spectroscopic techniques. NIRS has been discussed at depth earlier in this chapter. Continuous wave NIRS is perhaps the most appealing method from a standpoint of cost and technical difficulty, but is also the most prone to error due to modelling assumptions.

The wavelength dependence of the pulsatile PPG modulation depth forms the basis of pulse oximetry, which typically uses dual wavelengths; one in the far red (around 660nm to 730nm) and another in the near infra-red (around 860nm to 950nm). By measuring the modulation depth at two wavelengths it is possible to solve for both the magnitude of the pulsatile fluctuations (the PI discussed earlier) and the oxygen saturation of the pulsatile blood. As the pulsation is assumed to be purely arterial, this saturation is taken to be the arterial blood oxygen saturation. If the number of wavelengths in increased, a more complex set of equations can be solved, and the possibility may exist to obtain information concerning the static (i.e. non pulsatile) absorbers within the tissue.

A simple form of this concept was developed by Kohl et al [148]. Kohl used an off-the-shelf pulse oximeter to provide a measure of arterial saturation. The theoretical form of the PPG modulation depth D_{theory} versus wavelength assuming no path length

variation with wavelength was then calculated using published absorption data for Hb and HbO_2, and the normalized DPF was estimated using the normalized experimentally recorded modulation depth D_{meas} as $DPF = \frac{D_{meas}}{D_{theory}}$. Due to the unknown magnitude of the pulse induced arterial volume changes, only a normalized DPF can be obtained through use of this technique.

The differential path length factor technique gave credible results, but by using an off-the-shelf pulse oximeter Kohl was effective employing circular logic, since the pulse oximeter result is assumed to be accurate when the pulse oximeter is itself a two wavelength PPG modulation depth based device. In order to devise a modulation depth based fitting technique that does not rely on flawed logic, it will be necessary to either assume a typical arterial blood oxygen saturation of around 98%, or to include the arterial saturation into the model. As it involves making fewer assumption, it was decided to pursue the second route.

Although the differential path length factor NIRS technique is a simple and appealing one, it offers no insight into the cause of the path length variability. For this purpose a more thorough theoretical model of light transport is required. Patterson et al [186] present such a model; an application of diffusion transport to photon migration through scattering media. This diffusion model has been widely applied in NIRS tissue studies, where its assumptions of $\mu_s' \gg \mu_a$ and homogeneous tissue properties are most applicable.

For a more thorough analysis to be made, a model of the PPG induced by arterial pulsation is required. It is necessary to assume that the PPG is caused by a changing tissue blood fraction for this to be a simple method. Then, defining the PPG modulation depth as M:

$$M = \left(\frac{\partial I}{\partial \mu_a} \Delta \mu_a \right) / I_0, \tag{6.5}$$

$$\Delta \mu_a = \left(\mu_{a(blood)} - \mu_{a(tissue)} \right) \Delta V, \tag{6.6}$$

where ΔV is peak to peak change in tissue blood fraction over the PPG cycle, and

$$\left(\frac{\partial I}{\partial \mu_a} \right) / I_0 \tag{6.7}$$

is a function of our chosen model and our tissue properties ($\mu'_{s(tissue)}$ and $\mu_{a(tissue)}$). The form of 6.7 can be derived from the partial differential of equation 7 in the original

paper ([186]) by Patterson et al.

$$\left(\frac{\partial I}{\partial \mu_a}\right)/I_0 = \frac{-(z_0^2+\rho^2)}{2(D+\sqrt{D\mu_a(z_0^2+\rho^2)})}, \tag{6.8}$$

where ρ is the source to detector separation, $z_0 = 1/\mu_s'$, and D is the photon diffusion coefficient. Patterson assumed a diffusion constant of:

$$D = \frac{1}{3(\mu_a+\mu_s')}, \tag{6.9}$$

but there is controversy over the true form of D. Housley discusses discrepancies over the form of D in [187]. Elaloufi [188] discusses an equation of the form $D = 1/3(0.2\mu_a + \mu_s')$. It should be noted that the controversy over D mostly concerns the μ_a dependence, so in the near infrared case of μ_s' scattering dominance, the difference will be very modest.

Returning to equation 6.8, this can also be expressed as an effective path length from the modified Lambert-Beer law. From the modified Lambert-Beer law:

$$I = I_{initial} exp(-\mu_a \rho_{eff} + S), \tag{6.10}$$

differentiating 6.10 gives an alternative interpretation of 6.8:

$$\left(\frac{\partial I}{\partial \mu_a}\right)/I_0 = -\rho_{eff}, \tag{6.11}$$

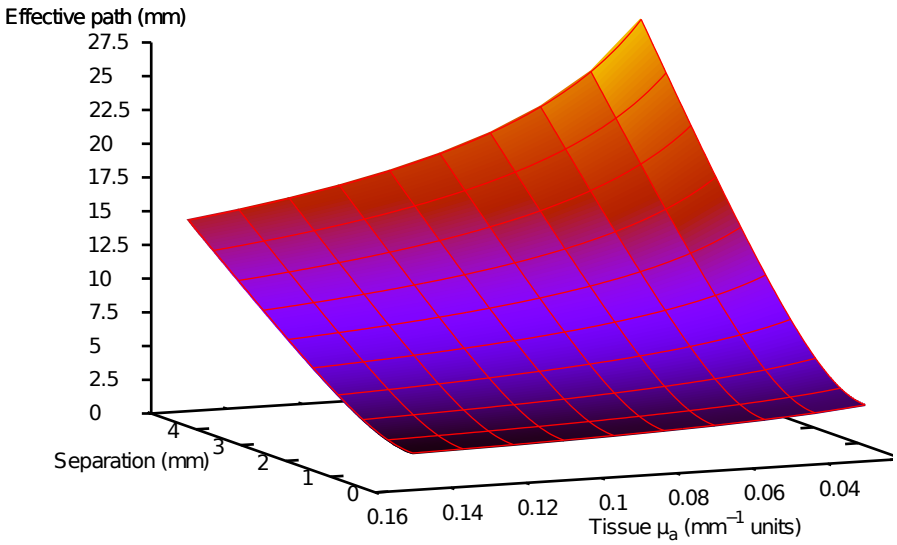


Figure 6.3: ρ_{eff} at 0 to 5mm source detector separation, $\mu_s' = 0.8mm^{-1}$

.

Note that $I_{initial}$ is not the same as I_0 introduced earlier, which was defined as the baseline intensity at the detector face.

This effective path length is plotted versus μ_a and ρ in figure 6.3. It can be seen that once the separation exceeds several mean transport free paths, the effective path is approximately linear with physical separation. However, the gradient of ρ_{eff} versus ρ is a nonlinear function of tissue absorption.

From equation 6.5, it can be seen that modulation depth will be proportional to the arterial blood absorption. If both the arterial blood and the fixed tissue absorption is assumed to arise from a mixture of Hb, HbO_2, and H_2O, then it can be seen that we have linear dependence on absorber concentrations in arterial blood and nonlinear dependence on fixed tissue absorbers. This linear versus nonlinear relationship suggests that with modulation depth measurements at a sufficient number of independent wavelengths, it may be possible to distinguish the effects of tissue absorption.

Equation 6.8 has a relatively complex form with a highly nonlinear dependence on both μ_a and μ_s', but as suggested by figure 6.3, approximations are valid at larger separations ($\rho \geq\sim 2mm$):

$$\left(\frac{\partial I}{\partial \mu_a}\right)/I_0 = \frac{-\rho^2}{2(D + \rho\sqrt{D\mu_a})}. \tag{6.12}$$

If a further assumption is made, that $\mu_s' \gg \mu_a$, this can be further simplified to:

$$\left(\frac{\partial I}{\partial \mu_a}\right)/I_0 = \frac{-\rho^2}{2(1/3\mu_s' + \rho\sqrt{\mu_a/3\mu_s'})}. \tag{6.13}$$

In typical tissues, with $\rho > 2mm$, $1/3\mu_s' < \rho\sqrt{\mu_a/3\mu_s'}$, and this simplifies further to:

$$\left(\frac{\partial I}{\partial \mu_a}\right)/I_0 = \frac{-\rho\sqrt{3\mu_s'/\mu_a}}{2}. \tag{6.14}$$

This is in agreement with an alternative form of the tissue attenuation derived from the diffusion model by numerous authors (for example Patterson and Flock [189]):

$$\mu_{eff} = \sqrt{3\mu_a(\mu_a + \mu_s')}, \tag{6.15}$$

as if the assumption $\mu_s' \gg \mu_a$ is made, equation 6.15 leads to the same form for 6.14.

Substituting 6.14 back into 6.5, and applying 6.6:

$$M = -\rho\sqrt{3\mu_s'/\mu_{a(tissue)}}\left(\mu_{a(blood)} - \mu_{a(tissue)}\right)\frac{\Delta V}{2}. \tag{6.16}$$

If a power law for is assumed for μ_s', i.e. $\mu_s'(\lambda) = k\lambda^{-h}$ then this equation can be rearranged into the form:

$$M = -\rho\sqrt{3k}\lambda^{-h/2}\left(\frac{\mu_{a(blood)} - \mu_{a(tissue)}}{\sqrt{\mu_{a(tissue)}}}\right)\frac{\Delta V}{2}. \tag{6.17}$$

Introducing a change of variable, $\Delta V' = \Delta V \rho\sqrt{3k}/2$,

$$M = -\lambda^{-h/2}\left(\frac{\mu_{a(blood)} - \mu_{a(tissue)}}{\sqrt{\mu_{a(tissue)}}}\right)\Delta V', \tag{6.18}$$

$$M\lambda^{h/2}/\Delta V' = \left(\frac{\mu_{a(tissue)} - \mu_{a(blood)}}{\sqrt{\mu_{a(tissue)}}}\right). \tag{6.19}$$

The volumetric tissue absorption is much lower than blood absorption, so this can be further approximated as:

$$M\lambda^{h/2}/\Delta V' = \left(\frac{-\mu_{a(blood)}}{\sqrt{\mu_{a(tissue)}}}\right). \tag{6.20}$$

The nonlinear function of tissue and pulsatile blood absorption on the right hand side of this expression can be found in arbitrary units by scaling the measured modulation depth M with wavelength $\lambda^{h/2}$, where h is taken from published results. Values for h of approximately 0.7 have been suggested by many authors [190].

If the PPG modulation depth is measured at wavelengths where there is high absorption from water and lipids within the tissue, it may be possible to find tissue absorption in absolute units, as the tissue absorption $\mu_{a(tissue)}$ will have water and lipid components that together account for approaching 100% of the tissue volume.

In order to test this modulation depth concept, near infrared PPG modulation depth data for the adult finger from Crowe et al [191] were used together with a numerical fitter using the fitting function sqp from the GNU-Octave programming language. Assuming $h = 0.7$ in line with published scattering data, the published modulation depth datasets were scaled with wavelength and normalized. The pulsatile blood was assumed to have a total Hb of 4.5mMol, in line with typical healthy adults, with the pulsatile water being assumed to have an insignificantly small effect. The arterial oxygen saturation was initially set at 97.5%, but was unconstrained in the fitter. Tissue absorption was assumed to arise from a mixture of lipids and water, summing to 100% of the total volume, with Hb and HbO_2 also present. This resulted in a model with four degrees of freedom, compared to 20 or more measurements in the raw data. Thus the fitter is over-constrained by the data, and the fitting process should be capable of producing stable solutions.

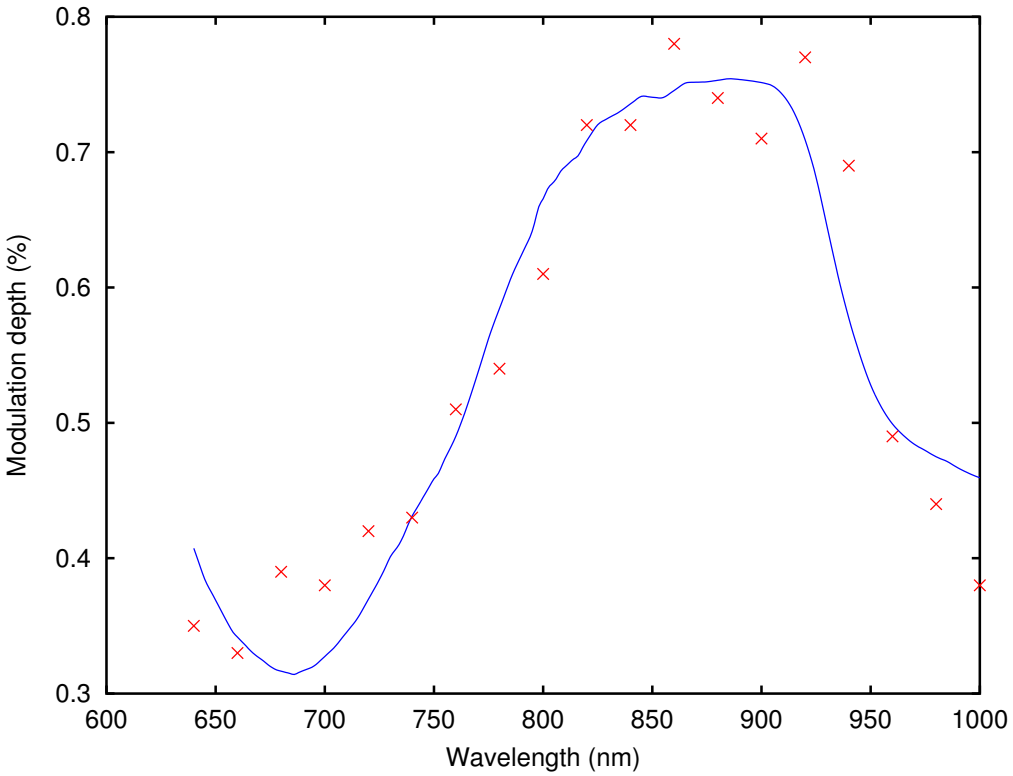


Figure 6.4: Reflected NIR modulation depth data from Crowe et al, subject 1

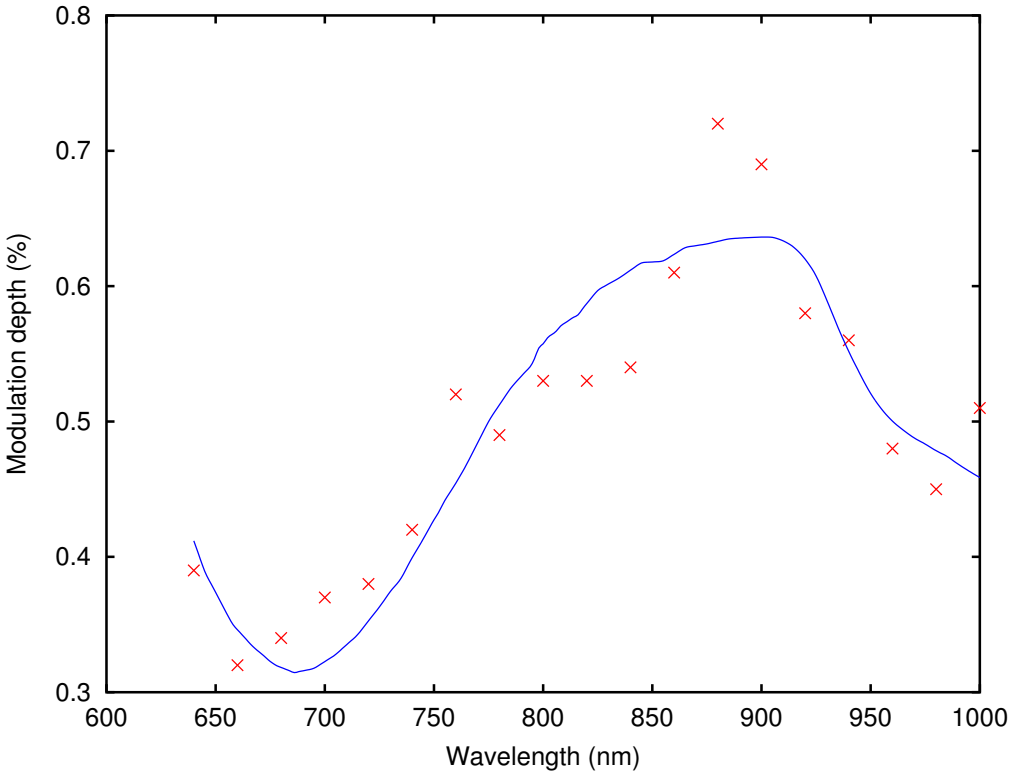


Figure 6.5: Reflected NIR modulation depth data from Crowe et al, subject 2

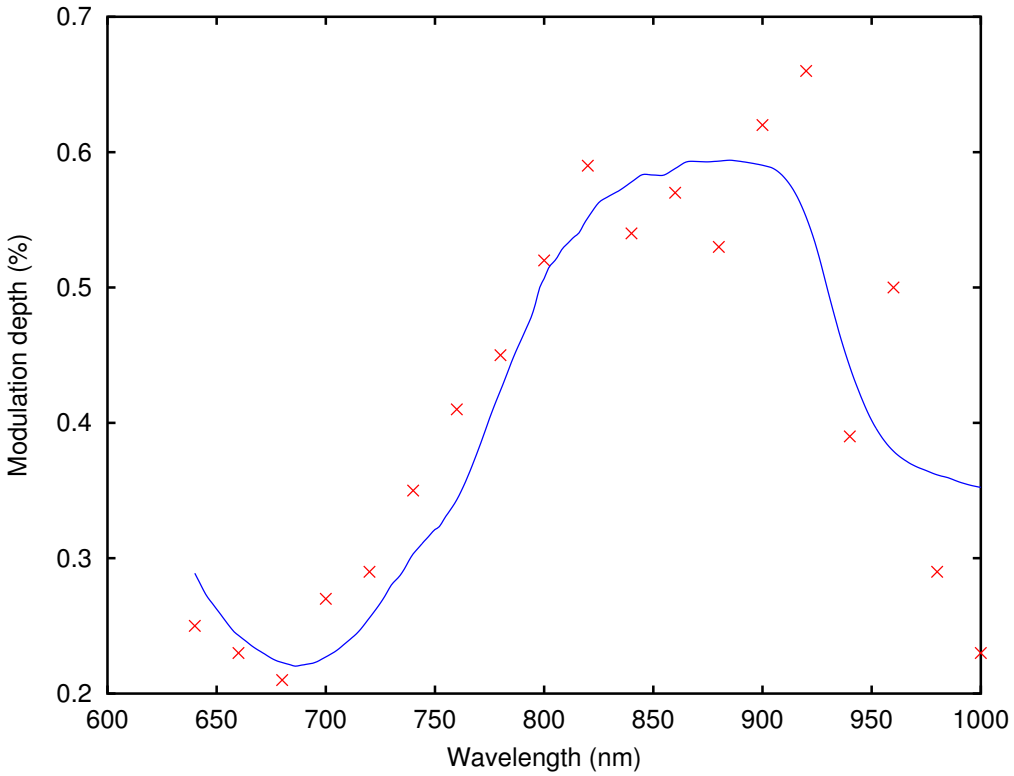


Figure 6.6: Reflected NIR modulation depth data from Crowe et al, subject 3

Subject	Arterial saturation	Tissue saturation	Tissue total Hb	Lipid (% vol)
1	94%	64%	0.254mMol	0%
2	95%	82%	0.564mMol	0%
3	90%	31%	0.233mMol	0%
3(exc)	97%	61%	0.178mMol	0%

Table 6.1: Fitting results using reflectance data from Crowe et al, and equation 6.20

Figures 6.4 to 6.6 show the results of fitting the equation 6.20 to the reflection data presented by Crowe et al. Table 6.1 gives the best fit tissue parameters and arterial blood saturations for the three subjects. All percentage values are rounded to the nearest percent, and the tissue molarity is given to 3 significant figures.

It can be seen that the arterial blood saturation is apparently significantly lower than the typical value of 97.5% found in healthy adults. Considering for example subject 3, it is also apparent that there are significant errors on the published modulation depth data. The modulation depth changes rapidly between 920nm, 940nm and 960nm, but considering the absorbers present in typical tissue, there is no clear cause of this change other than interference such as motion artefact or fluctuations in the magnitude of the volumetric pulsations. The equipment used to record the data, a white light

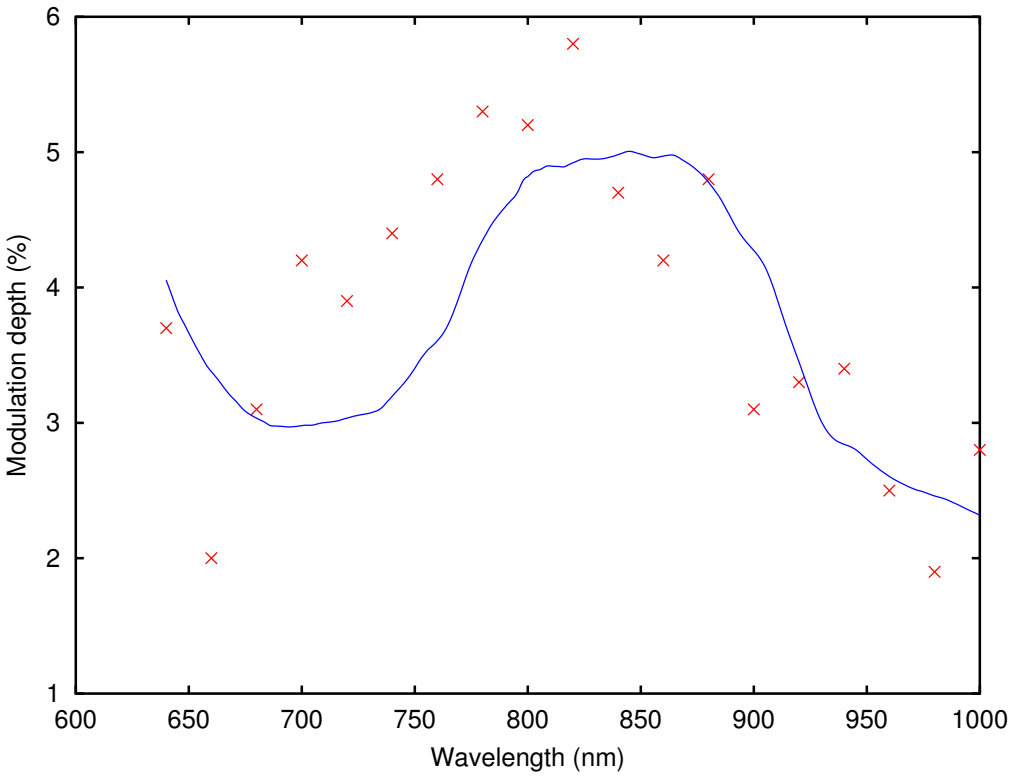


Figure 6.7: Transmitted NIR modulation depth data from Crowe et al, subject 1

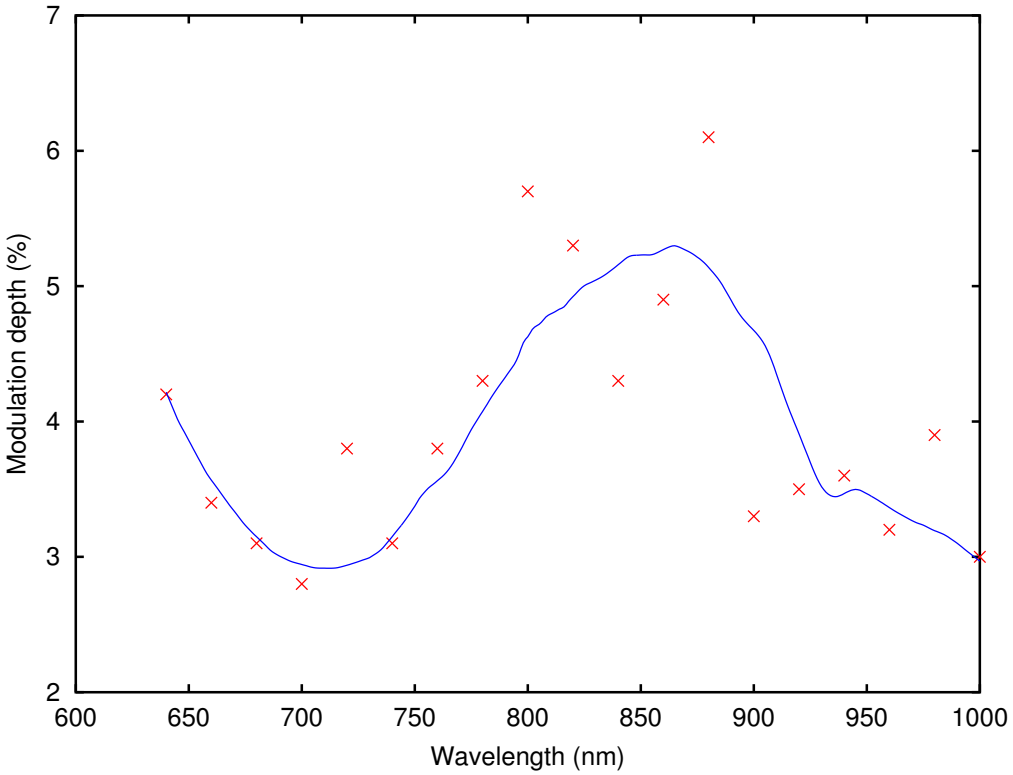


Figure 6.8: Transmitted NIR modulation depth data from Crowe et al, subject 2

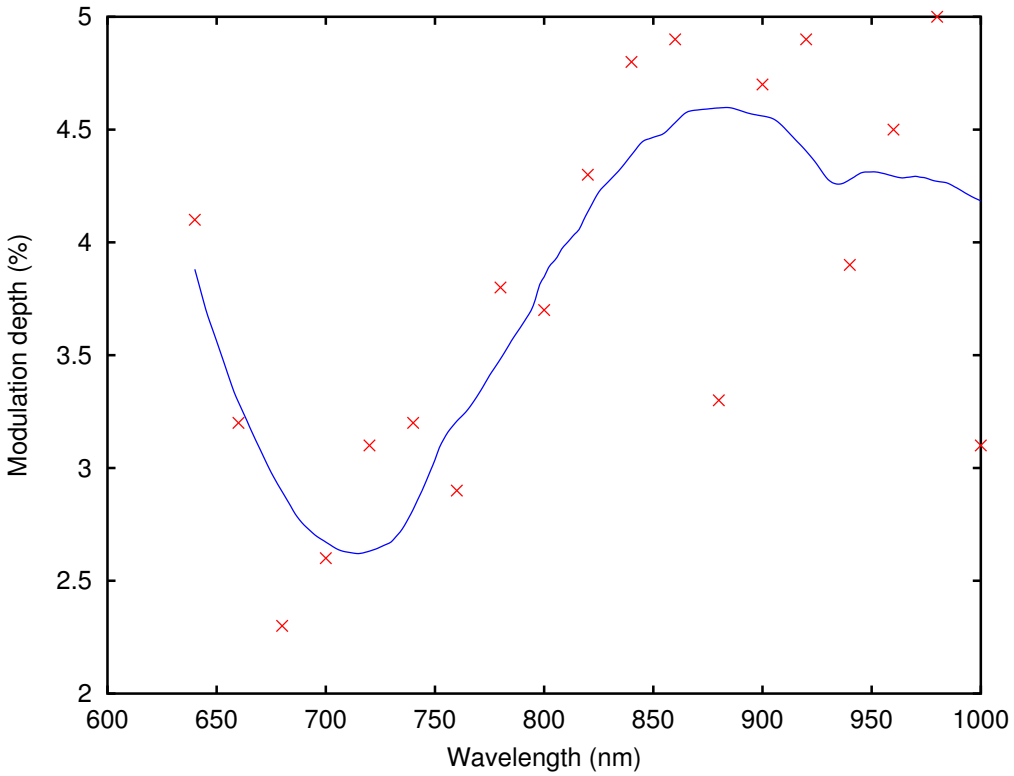


Figure 6.9: Transmitted NIR modulation depth data from Crowe et al, subject 3

Subject	Arterial saturation	Tissue saturation	Tissue total Hb	Lipid (% vol)
1	94%	72%	0.036mMol	58%
2	75%	12%	0.084mMol	70%
3	71%	19%	0.737mMol	69%

Table 6.2: Transmission results using data from Crowe et al, and equation 6.20

source followed by a stepper motor controlled monochromator, leads to a long acquisition time as the wavelength was switched through 20nm increments. It can be seen that this likely resulted in variations in the PPG amplitude caused by factors such as the respiratory cycle aliasing with the wavelength steps and causing rippling errors on the PPG amplitude data set.

Assuming a blood haemoglobin molarity of around 5mMol (typical for healthy adults), then the total tissue haemoglobin molarities are consistent with a blood volume fraction of between 5 and 10%. This is in agreement with published figures for the adult finger. The tissue saturation for the first two subjects is also in agreement with published values, but the fit to subject 3 has an excessively low saturation. The subject 3 fit was repeated with the data at wavelengths greater than 900nm excluded in an effort to reduce the impact of the apparently erroneous points in the data. This resulted in fitted tissue

properties in line with published data for healthy adults, as can be seen from the last line of the table.

The equivalent transmission data are presented in figures 6.7 to 6.9. As evident from the fit parameters in table 6.2, the transmission data did not result in credible tissue parameters.

It should be noted that the reflection data were recorded using a path length through the tissue of around 5mm, i.e. not diffuse reflectance. The transmission data were recorded through an adult thumb, giving a path length of around 15mm. Both of these separations are sufficiently large that equation 6.20 should be valid. However, the data quality from the transmission recording is much worse, with a very significant scattering of the modulation depth measurements evident on the graph, especially at the longer wavelengths. The reason for this is unclear, but it may be partly due to the much lower light intensity that would have resulted from the longer path length through tissue in the case of transmission.

Damianou gives further details of the experimental procedure in his thesis [192], figure 3.3 shows the hardware design, and section 3.7.1 gives a description of the signal processing algorithm used to calculate the modulation depth or "AC/DC". Digitised AC filtered data were passed through a rolling maxima-minima window. The averaged output value was assumed to be equal to the PPG amplitude. The baseline or "DC" was taken directly through an analogue low pass filter and a second ADC channel. This algorithm would appear to be vulnerable to artefacts falling within the same frequency spectrum as the PPG, which would pass through the analogue bandpass filter, causing minima and maxima to be distorted. An adaptive notch filter tracking the heartrate may have resulted in improved results from the experiment.

6.7 Capillary refill

The capillary refill test measures the time taken for blood to refill the capillary bed after it has been emptied by pressure applied from the surface. The test is widely used clinically in a manual form: the clinician will apply firm pressure with a finger for 5 seconds, before releasing and manually assessing the time taken for the blanched mark to fade until the colour matches surrounding skin.

There has been discussion of the errors inherent in the purely manual method, for example Lobos et al [193] found widespread variation in not only the interpretation of capillary refill time, but also the correct application pressure and the time for which pressure should be applied. Brown found adverse lighting conditions lead to bias and

high variability in manually assessed refill time [194].

Capillary refill is clearly a feasible target for automation, which could potentially lead to dramatic improvements in consistency of both pressure application and measurement of the blanching. Automation of capillary refill could take the form of either manual pressure application with opto-electronic measurement of the refilling blood (for example NIRS), or automated pressure application (for example pneumatically). A combination of both automated pressure application and optical sensing could lead to a fully automated device, an innovation that is not currently marketed by any medical device manufacturer, although US Patent 8082017 [195] describes a modified finger probe pulse-oximeter used to record capillary refill. A similar device was demonstrated by Kviesis [196], and video processing in conjunction with manual pressure application has been demonstrated by multiple authors, for example Shavit [197]. It should be noted that the effective path length of the observed light is much shorter in video studies than is usually the case in NIRS, due both to the fact that most video studies have made use of diffuse reflectance measurements, rather than the physically separate source and detector arrangement common in NIRS, and also due to the use of shorter optical wavelengths in the visible range, meaning both μ_s' and μ_a are larger than in the near infrared window region.

6.7.1 Existing clinical use

The earliest clinical use of capillary refill time was as an indicator of dehydration, suggested as early as 1910 by Takayesu et al [198], but the authors make it clear that the test was already commonly known by that time. By 1940, Guedel et al [199] were using capillary refill time measured on the forehead during surgery as an indicator of surgical shock. In 1947, Crismon et al [200] used the capillary refill effect to study the effects of frostbite in rabbits. Since the 1940s, capillary refill has become more widely used as a simple indicator for shock and dehydration.

Use of capillary refill as part of a numerical trauma score was first popularized by Champion in 1981 [201]. Champion defined 2 seconds as a threshold time, with a capillary refill of more than 2 seconds being indicative of shock. However, Baraff was highly critical of Champion's work [202], and noted that there had been few comprehensive studies into both the nature of the capillary refill signal and its variability amongst both healthy patients and patients suffering from conditions such as shock or dehydration. He also states that it was not proven that capillary refill is a direct measurement of perfusion. Schringer reached similar conclusion regarding use of a simple cut-off time for defining normal capillary refill [203], with age, sex and temperature leading to sig-

nificant changes in capillary refill time in healthy volunteers. However, Gorelick [204] found an approximate linear relationship between temperature and capillary refill time, suggesting that normalizing measured capillary refill times to an equivalent time at a defined temperature (25°C for example) would remove a significant proportion of the variability. In a 1989 revision to his trauma score [205] Champion removed capillary refill time on the grounds that "it was difficult to assess in the field". Saavedra et al [206] found refill time to be a good indicator of hydration state, suggesting that hydration state may be a confounding factor complicating any attempt to use capillary refill as a measure of perfusion.

Since the 1980s, capillary refill has been widely touted as a direct measure of perfusion, for example by Lima and Bakker [83]. However, there are relatively few published studies directly comparing capillary refill with a gold standard measure of perfusion. Osbourn [13] found that capillary refill time was a poor indicator of low peripheral flow in paediatric care. Figure 2 from Osbourn et al shows a scatter plot of ultrasonically measured blood flow versus capillary refill time. It can be seen that a long capillary refill time has good specificity for perfusion, but not good sensitivity. However, use of trend data to detect changes in perfusion was not studied by Osbourn. Gray et al [207] found no significant correlation between capillary refill time and perfusion measured via Doppler ultrasound.

It is clear there is some confusion regarding the nature and applicability of the capillary refill time to perfusion. It is interesting to compare capillary refill to the other methods of perfusion assessment that have been considered so far, all of which have their origin more in the applied sciences than the clinical sphere. Capillary refill time in contrast has been used almost exclusively as a manually applied phenomenological metric, with no underlying theoretical backing.

6.7.2 Hypotheses concerning the origin of the refill signal

As Baraff commented in 1993 [202], there has been little development of any theoretical explanation of capillary refill time. It is widely assumed that once the capillary bed has been emptied of blood, arterial flow is the sole source of refilling blood. However, conclusive evidence is lacking. In 2011 Pickard could find no published studies demonstrating a clear basis for claims that capillary refill is directly related to perfusion [208].

It may be of interest to compare capillary refill to the controversy over bioheat equations. Pennes's original bioheat equation is a simple model where small clusters of capillary loops are perfused from an arterial supply. This could be compared to a simple capillary refill model with all refilling blood coming from the arterial side. In reality the Pennes

model was found to be somewhat inaccurate, partly as the micro-vascular system is very complex and highly interconnected. Applying this finding to capillary refill, it may be the case that refilling blood originates partly or mainly from the surrounding capillary bed, and is able to flow into the blanched area due to the interconnections within the capillary bed.

It may further be noted that when pressure is released from the skin surface, the capillary bed at the edge of the blanched zone is often observed to refill first, and never the converse. This observation would be expected if the blood originates primary from the surrounding capillary bed, as due to the relatively high flow velocities found in the arteries and arterioles, a simple arterial refill model would give approximately even refill time across the blanched area.

In a 2012 study, Lima et al investigated the effects of entire body cooling on peripheral circulation [209], using a Masimo pulse oximeter with PI output, a NIRS device applied to the Thenar Eminence combined with arterial occlusion, temperature probes, and manually assessed capillary refill time (CRT) in eight healthy adults. A statistically significant CRT change was seen upon cooling. A InSpectra model 650 NIRS probe measured tissue saturation and total haemoglobin throughout the arterial occlusions and recovery. It should be noted that, due to the design of the experiment, there was little if any direct cooling of the hand or forearm, as the instrumented arm was placed outside the cooling blanket used to expose the body to low temperature, so therefore it is unclear if the results from Lima et al are directly comparable to water bath cooling of the hand and forearm.

Lima et al found that no significant change in total haemoglobin with whole body cooling, with the authors concluding that temperature induced changes mostly affected the arterial compartment, and as the NIRS device was most sensitive to the capillary and venous compartments, it had poor sensitivity to arterial vaso-constriction. As the Lima study did not use venous occlusion, it is not possible to directly compare their results to studies of volumetric perfusion. However, the fact that Lima saw no significant change in the rate of StO_2 decrease during arterial occlusion with reduced temperature suggests that the assumption that there is no significant change in metabolic uptake of oxygen when the limb is cooled is valid, and that therefore the Fick principle can be applied to oxygen extraction index.

The significant change in capillary refill time seen by Lima et al may be explained as arising via the vaso-constriction within the capillary bed rather than as a direct consequence of reduced perfusion, but the fact that no significant change in total tissue Hb was seen would suggest that there was no significant vaso-constriction within the

capillary bed, and that therefore it might be reasonable to conclude that perfusion played a role in the capillary refill time change.

6.8 Conclusion

A review of optical techniques of relevance to non-invasive peripheral perfusion monitoring has identified a significant number of promising methods.

Laser Doppler is perhaps the classic optical perfusion technique, but a review of relevant literature and theory leads to the conclusion that it cannot provide an absolute measure of perfusion. It would appear that a measure of perfusion changes, i.e. perfusion trend, can be reliably obtained using this technique. Use of more sophisticated techniques such as DWS and DCS in place of simpler QELS techniques theoretically offers the potential of volumetric perfusion, but published data indicate that the complex motion exhibited by scattering red blood cells during their passage through the capillary bed results in little gain in the applicability of the available information. Laser Doppler systems are available off-the-shelf, and the technique has seen relatively widespread adoption within the research community, as well as some limited use in routine clinical monitoring, so although it clearly offers some promise as a means of non-invasive perfusion monitoring, laser Doppler does not seem a promising candidate technology for meaningful further development.

NIRS technology has reached a similar readiness level to laser Doppler, with off-the-shelf devices on the market, but seeing limited use in the clinical environment. There is a cost/performance tradeoff between different NIRS techniques, with continuous wave having significantly lowest cost, but also the worst performance. For an approximate metric of perfusion adequacy, continuous wave may be a suitable low cost solution. However, the suitability of tissue oxygenation as a reliable indicator of compromised perfusion is not completely proven, so more "active" techniques such as venous occlusion might be used in conjunction with NIRS. A large number of studies have demonstrated use of venous occlusion with NIRS to measure volumetric perfusion. A novel technique involving use of PPG modulation depth to obtain tissue oxygen saturation was attempted, and yielded promising results, although the quality of the raw data is in doubt. A patented algorithm for processing pulse oximeter data was evaluated using numerical simulation and appears to offer the potential of measuring venous saturation, potentially a much more useful metric of perfusion adequacy than overall tissue saturation.

Capillary refill time is widely used clinically as a manual technique. However, the evidence for this metric as a measure of perfusion would appear to be lacking, the 2012 publication by Lima et al [209] being the only clear published work that could be

found to support the use of capillary refill as providing any indication of true volumetric perfusion, rather than merely as a measure of vaso-constriction and blood viscosity. It is somewhat perplexing that there has been so little development work towards automation of capillary refill, as this would seem relatively easy from a technical point of view. From the few published studies, it would seem sensible to conclude that capillary refill time is unlikely to provide a direct measurement of perfusion, but the precise nature of the signal is unclear, and detailed studies may well provide useful additions to the state of understanding.

In addition to the capillary refill, Lima at al also found significant changes in the PI, and core-peripheral temperature, apparently demonstrating the suitability of all three techniques for perfusion monitoring. It is notable that PI would appear to be a more effective indicator of the temperature induced reduced perfusion state than either capillary refill time or differential temperature, which both show more significant fractional inter subject variation, and CRT outperforming differential temperature by a small margin. However, a study by Zaramella et al [180] found that although correlated with perfusion, PI was a relatively inaccurate gauge of volumetric perfusion, at least in healthy neonatal subjects. Looking at inter subject variation in study of cooling induced changes by Lima et al, CRT outperformed differential temperature by a more considerable margin.

Alongside laser Doppler and NIRS, Capillaroscopy is a third commonly used off the shelf technology. It does not appear well suited to continuous monitoring, being bulky and expensive in its current form. It may be feasible to dramatically reduce the size and cost, but use on thick skin leads to high scattering, and current devices struggle to obtain usable images on all but the thinnest skin types. Use of capillaroscopy in conjunction with coherence tomography may overcome some of these problems, but this comes at the expense of dramatic increases in complication and likely cost of any device.

Chapter 7

Choosing and developing a candidate technique

With a thorough analysis of potential non-invasive peripheral perfusion monitoring techniques having been conducted, it should now be possible to evaluate their suitability in a systematic manner, and pick a candidate technique or set of technologies that would be best to develop further.

This is not a simple process, as there is a complex set of trade-offs to be made in meeting the goals. First it would be prudent to consider what the aims are for this research. In the introductory chapter, a goal was set to develop a non-invasive monitoring solution for the peripheral perfusion, with a particular focus on reliable identification of flow shutdown, or decrease of flow to dangerously low levels. It was also noted that for any device to become successful within the clinical environment, cost and ease of use are of paramount importance. Note for example the success of pulse oximetry, which is so cheap to implement that the sensors are disposable in many cases, and can be clipped onto a finger or adhered to the wrist or ankle by staff who have received minimal training. For this research to lead to anything other than a research device, or worse a laboratory "curiosity", these considerations of real world practicality must be made.

In order to accurately choose a monitoring method from the large number of candidates that have been identified, it would be beneficial to employ some form of scientific decision making process. Operations research is an area of management science that focusses on such problems. Specifically, multiple-criteria decision analysis or MCDA addresses multi-parametric decision making based on trade-offs. The Analytic Hierarchy Process (AHP) developed by Thomas L. Saaty [210] is perhaps the most commonly employed MCDA algorithm. Saaty later developed AHP into a more generalized form; the An-

alytic Network Process (ANP) [211]. ANP extends the simple weighted hierarchical decision process used in AHP to allow more complex interrelationships between choices. In the ANP process, each decision criterion is a node in a network, allowing connections between nodes to transfer weighting. As it is not immediately apparent how the complex network system of ANP could be applied to the decision making process, it was decided to proceed using AHP, at least for an initial assessment.

7.1 Analytic Hierarchy Process (AHP) based analysis

AHP assigns weights to alternative options based on a series of weighted criterion. In industry, AHP is typically used as part of a process such as BORC (Benefits, Opportunities, Risks, Costs). Where the AHP is often used for one or more of these four parameters in isolation, before the scores are combined using an equation either of the form BO/CR or $B + O - C - R$. Not all the methods discussed previously are considered here; NMR was judged to be impractical, USMOT was not found to provide useful information for perfusion assessment, and it is unclear if oxygen extraction index derived via DST can yield useful perfusion information.

For the purpose of this work, opportunities is not particularly relevant, as a business case is outside the scope of this work. Cost could be defined as the resources in terms of time and money required to bring the project to fruition. Risk could be defined as the probability that the research and development effort will fail due to either overspending or overrunning its predicted costs and duration, or failing to result in any solution to the problem. It may also be feasible to relate risk to the current Technology Readiness Level (TRL) of the chosen technology. Benefits are more difficult to quantify, but it would appear that benefits evaluation is where AHP could be practically applied.

A set of evaluation criterion must be defined before applying AHP. To be able to draw up this list of benefit criterion, the definition of benefits must be clarified. To do this, the goal of the research and development effort must be clearly established in terms of a deliverable target. Here, the deliverable is not a fully developed medical device, but the core technologies behind such a device, which can be transferred in the form of intellectual property and/or patents to a medical device manufacturer. Thus it is appropriate to separate the core technology research and development costs and risks from the manufacturer's costs (both in terms of one-off setup and production costs per finished device) and risks (the MRL or Manufacturing Readiness Level can be applied here).

The goals discussed in the introduction to this chapter can also be considered, leading

finally to a list of six benefit criterion for use in the AHP assessment.

1. The likely quality of the perfusion metric obtained from the candidate technology.
2. The practicality of application and use of the technology in a realistic clinical environment.
3. The risk to the patient or clinician posed by the technology; considering worst case failure mode.
4. Likely cost of such a device when put into mass production (i.e. operating costs/patient).
5. Manufacturing Readiness Level (MRL).
6. One-off manufacturing costs and difficulty; proportion of off-the-shelf components used etc.

As this is fundamentally a research oriented project, the weighting assigned to the first two criterion will be disproportionally high relative to for example a similar project in industry. It is also notable that risk can often be mitigated at the expense of a more complex and thus expensive device. Thus it can be seen the ANP may have a role to play in analysis of this problem, as for a device to enter market its risk must have been reduced to acceptable levels, creating a linkage between risk and the three manufacturing related criteria. Furthermore, medical devices would ideally have a risk level that is sufficiently low to be irrelevant, and so risk could be eliminated from the final output entirely, and serve only to reduce the three manufacturing related scores through network linkages.

It is interesting to consider commercial transcutaneous blood gas sensors at this point. These are now used relatively rarely compared to pulse-oximetry. Transcutaneous blood gas sensing involves heating the skin to temperatures of around 40°C, an inherently risky procedure. The need to mitigate risk to the patient with very accurate temperature control compounded with the existing manufacturing challenges in producing accurate gas sensitive electrodes, resulting in a complex feedback system between criterion 2 to 6 in the earlier list. Looking at the situation using ANP, it can be seen that this has effectively caused a "vicious cycle" of negative feedback effects, with the current generation of transcutaneous blood oxygen and carbon dioxide sensors being neither particularly cheap nor practical or easy to use, and thus losing out to pulse oximetry in most clinical applications.

Unfortunately it is unclear how realistic node linkages could be introduced into ANP without much experience in the medical device field. Simple attempts to guess at likely

Criteria	6	5	4	3	2
1	4	3	2	2	1.5
2	4	3	2	2	
3	3	2	2		
4	2	2			
5	1.5				

Table 7.1: The AHP relative weightings assigned to the six criteria.

linkages would probably result in highly unrealistic outputs due to the nonlinear behaviour of networks with feedback loops. For this reason it was decided to proceed with a simple AHP method, but to assign scores based on an understanding of the likely impacts across the criteria. For example, a high risk would lead to a strongly negative benefit due not to risk being present in the production device, but due to the knock on consequences of risk mitigation.

As described by Saaty, the first step in an AHP analysis is to assign relative weights to the criterion using a correlation triangle. Table 7.1 shows the assigned relative weights.

The weights are relative assignments of importance of the row criteria over the column criteria. For example, criterion 1 is 4 times as important as criterion 6. There is clearly scope for some ambiguity here, with comparisons between the relative importances leading to logic contradictions. Saaty however designed the assignment technique to be as simple to use as possible, with logical contradictions allowed at the initial stage, only to be resolved using matrix algebra at a later stage. It can be seen that the chosen weighting is relatively simple, with the high numbered criteria being of lower importance than the lower numbers.

Thus to summarise the table:

- We are not too concerned over one off manufacturing costs - medical devices are a high volume market, and tooling costs will be relatively insignificant.
- A less practical but better performing device is only slightly preferable to a practical but poorly functioning solution. A similar but stronger preference exists for a more costly but better performing solution.
- We are relatively tolerant of manufacturing risk.
- Finding a solution with a minimal potential for patient harm (at the initial pre-commercialisation stage) is around half as important as finding an effective solution.

With a populated weighting table, the next stage of the AHP algorithm [212] is to create a matrix with the upper and lower triangular parts being equal to the corresponding elements from the weighting table and their reciprocals respectively. The diagonal is set to all ones. Using table 7.1 we arrive at the following form for the AHP weighting matrix:

$$W_{AHP} = \begin{pmatrix} 1 & 1.5 & 2 & 2 & 3 & 4 \\ 0.67 & 1 & 2 & 2 & 3 & 4 \\ 0.5 & 0.5 & 1 & 2 & 2 & 3 \\ 0.5 & 0.5 & 0.5 & 1 & 2 & 2 \\ 0.33 & 0.33 & 0.5 & 0.5 & 1 & 1.5 \\ 0.25 & 0.25 & 0.33 & 0.5 & 0.67 & 1 \end{pmatrix} \tag{7.1}$$

The eigenvector of this matrix is then found (to 2s.f. and where *eig* is a robust eigenvector approximation as described in Saaty's 1990 paper):

$$eig(W_{AHP}) = \begin{pmatrix} 0.29 \\ 0.27 \\ 0.17 \\ 0.13 \\ 0.08 \\ 0.06 \end{pmatrix} \tag{7.2}$$

Next, a matrix of "performance metrics" is created using both data and/or assessments of the perfusion data quality, device practicality etc. Clearly this is not trivial, it is necessary to make an assessment and assign a numerical score to all the criteria for each candidate technology. To simplify matters it was decided to assign numerical scores in the range of 0 to 5 to all the categories except the MRL, where the reciprocal of the MRL was used as a measurement of risk. This results in a non-linear relationship between MRL and risk score, which is likely to be more realistic, as failure is far more likely to occur in the early stages of most projects. The scores are shown in table 7.2. Note that a negative "benefit" is assigned to the danger, risk, cost, and set up difficult.

After normalizing columns, as described by Saaty, the following matrix is formed (7.3). The fifth criterion, MRL, is somewhat difficult to define at the start of the project. The score in the table is based on the extent to which off-the-shelf parts and standard manufacturing techniques can be employed. For example, Active Thermal sensors can be constructed very easily using polyamide based flexible PCB material coated with bio-compatible acrylic pressure sensitive adhesive. Surface mount thermistors can be used for both heating and temperature sensing. It can thus be seen that there is the

Technology	Criterion						TRL	"Cost"
	1	2	3	4	5	6		
Active thermal	4	4	-4	-2	7	-1	4	2
Passive thermal	3	3	-1	-3	4	-2	3	3
Doppler ultrasound	3	3	-3	-4	9	-3	7	4
Photo-acoustic	4	2	-4	-5	6	-5	6	5
QELS	2	3	-3	-4	10	-2	9	4
DWS/DCS	2	4	-3	-4	9	-3	7	5
NIRS(CW)	3	4	-2	-3	10	-2	9	3
NIRS(CW) + venous occlusion	5	2	-3	-2	6	-3	7	3.5
Capillaroscopy	4	2	-2	-4	4	-4	9	4
PPG PI	3	4	-2	-1	10	-1	9	1
Capillary refill time	4	4	-3	-2	4	-4	3	4

Table 7.2: The AHP scores assigned to the six criteria. The "cost", defined as relative estimated effort (in range 1 to 5) required to develop the technology, and the TRL are also shown.

"Capability to produce systems, subsystems or components in a production representative environment" meeting the criterion for MRL 7 (as outlined by Morgan [213]). In contrast, passive thermal monitoring has an MRL of only 4, as accurate measurement of thermal flux would require more unconventional manufacturing, for example integration of peltier elements with flexible printed circuit boards, something that could probably only be defined as "Capability to produce the technology in a laboratory environment" or MRL 4.

$$W_{scores} = \begin{pmatrix} 0.108108 & 0.114286 & -0.133333 & -0.058824 & -0.081707 & -0.033333 \\ 0.081081 & 0.085714 & -0.033333 & -0.088235 & -0.142987 & -0.066667 \\ 0.081081 & 0.085714 & -0.100000 & -0.117647 & -0.063550 & -0.100000 \\ 0.108108 & 0.057143 & -0.133333 & -0.147059 & -0.095325 & -0.166667 \\ 0.054054 & 0.085714 & -0.100000 & -0.117647 & -0.057195 & -0.066667 \\ 0.054054 & 0.114286 & -0.100000 & -0.117647 & -0.063550 & -0.100000 \\ 0.081081 & 0.114286 & -0.066667 & -0.088235 & -0.057195 & -0.066667 \\ 0.135135 & 0.057143 & -0.100000 & -0.058824 & -0.095325 & -0.100000 \\ 0.108108 & 0.057143 & -0.066667 & -0.117647 & -0.142987 & -0.133333 \\ 0.081081 & 0.114286 & -0.066667 & -0.029412 & -0.057195 & -0.033333 \\ 0.108108 & 0.114286 & -0.100000 & -0.058824 & -0.142987 & -0.133333 \end{pmatrix} \tag{7.3}$$

This matrix is multiplied by the eigenvector to give a score for each candidate technology (equation 7.4) and then normalized so the scores sum to 1.

$$W_{scores}eig(W_{AHP}) = \begin{pmatrix} 0.102 \\ 0.195 \\ 0.117 \\ -0.105 \\ -0.017 \\ 0.026 \\ 0.192 \\ 0.136 \\ 0.006 \\ 0.272 \\ 0.151 \end{pmatrix} \tag{7.4}$$

The scores in 7.4 are "benefits". There are also "costs" and "risks" to be considered. As with the treatment of MRL, risk as defined as $1/TRL$.

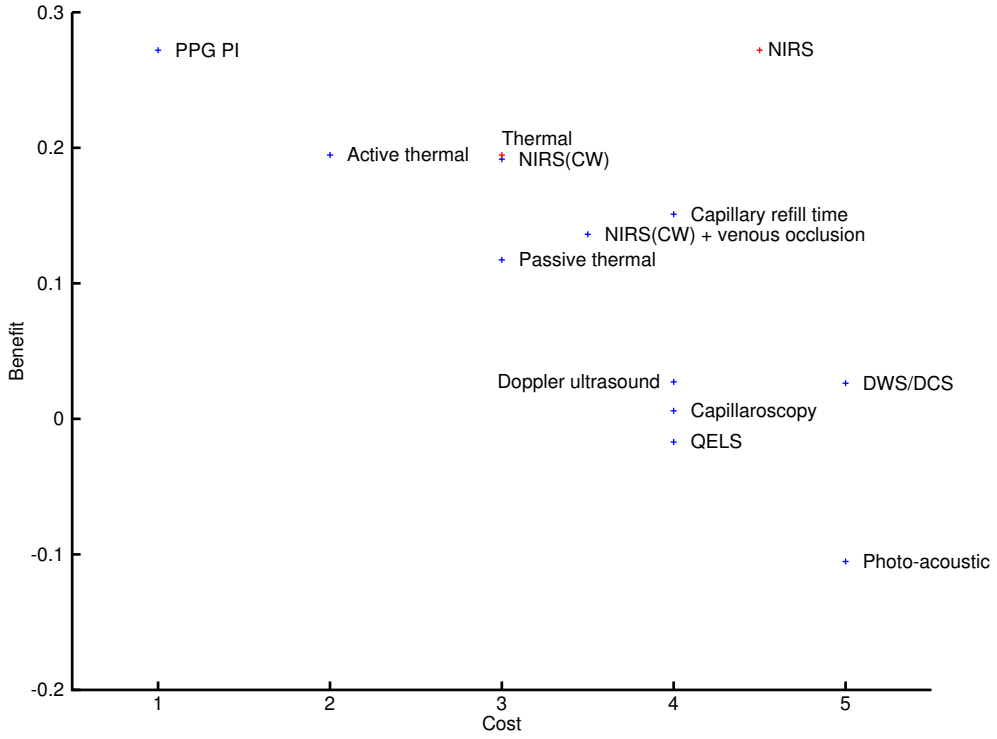


Figure 7.1: Benefit versus Cost for the candidate technologies

From figure 7.1 it can be seen that the candidate technologies are approximately aligned

on a line from low cost, high reward at the top left to low reward and high cost at the bottom right. This stands in contrast to the typical line passing close to the origin, but can be seen to be feasible as we have included manufacturing risk as a negative benefit, so complex expensive to implement techniques far from commercialisation readiness also tend to require the greatest development input. Off-the-shelf PPG based perfusion index (PI) is a clear winner. Active thermal monitoring is a relatively close second place, and NIRS techniques, passive thermal monitoring and capillary refill form a cluster of approximately equivalent evaluated suitability.

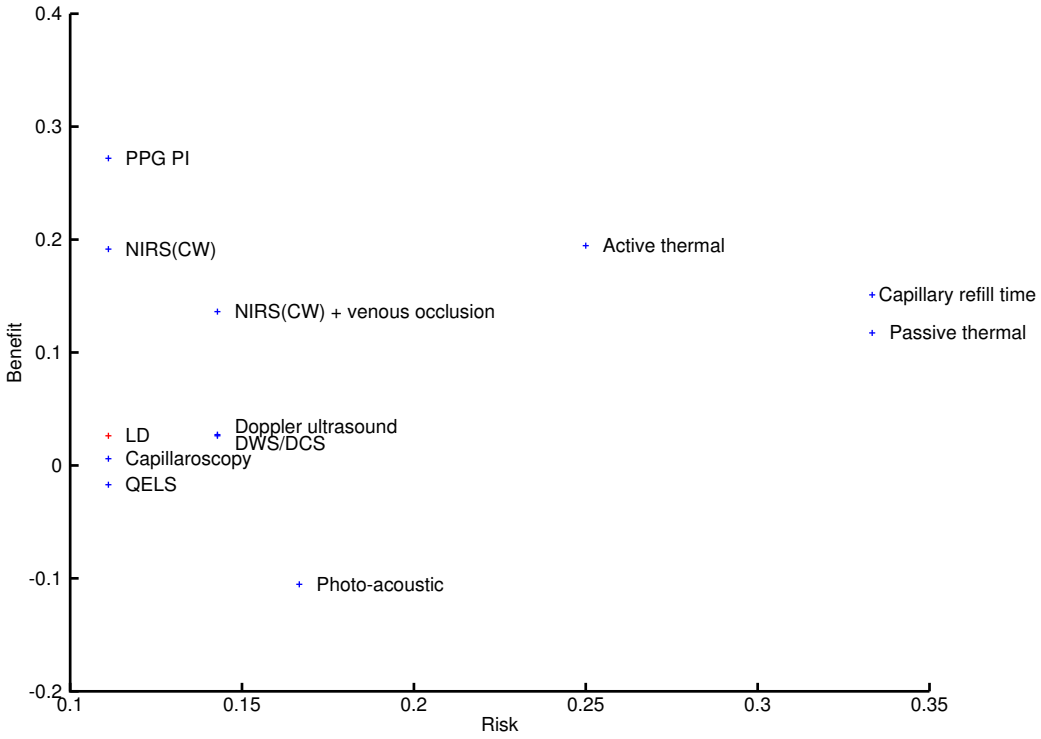


Figure 7.2: Benefit versus Risk for the candidate technologies. LD indicates the Laser Doppler cluster

In figure 7.2 it can be seen that there is more of a spread of risks. PPG PI is still a clear winner, but there is a large spread of risk amongst the previously clustered techniques below it. Continuous wave NIRS is a low risk alternative to PPG, and NIRS can be augmented with venous occlusion by addition of an off-the-shelf blood pressure cuff.

This suggests an alternative way to analyse the data; some technologies can be shared between multiple techniques. Thus some of the techniques can be "clustered" as they rely on common technology. The thermal techniques (active and passive) can clearly be clustered to a certain extent, as can the NIRS techniques (NIRS, NIRS + venous occlusion, capillary refill time and PPG PI). The Laser Doppler techniques DWS/DCS

and QELS can be clustered, but Doppler Ultrasound, Photo-acoustic and Capillaroscopy cannot be sensibly clustered in with other techniques. This clustering results in a dramatic reduction in risk, at only modest increase in cost. However the analysis is not trivial, for example consider the PPG PI technique; this has a low risk and cost as off-the-shelf equipment can be used. If we were to develop NIRS technology and apply to PPG PI, the risk would clearly be much higher than using an off the shelf PPG unit. Additionally, the risk is in arbitrary units, so a mathematically accurate estimate of the risk for a "cluster" is not feasible.

A simple method of clustering was attempted instead; the risk of a cluster was assumed to be equal to the smallest risk of its component techniques, the benefit equal to the highest component benefit, and the cost equal to the sum of "core" cost (the smallest cost out of the components) and the excess cost over the core technology incurred by each additional component. We thus have three clusters:

1. Laser Doppler cluster: QELS, DCS/DWS
2. Continuous wave NIRS cluster: NIRS, NIRS + venous occlusion, Capillary refill time, PPG PI (excluded from calculations)
3. Thermal cluster: Active and Passive.

In figure 7.2, the Laser Doppler cluster is indicated as "LD". The NIRS and Thermal clusters will be located at PPG PI and Active Thermal, and so have not been displayed. It can thus be clearly seen that development or acquisition of a continuous wave NIRS system that can be applied across the techniques in the second cluster offers the best reward to risk ratio.

Likewise for figure 7.1, the thermal and NIRS clusters are shown. The Laser Doppler cluster will be at the same location as DWS/DCS, and so has not been shown. Thermal is the best choice on a Cost/Benefit basis, but NIRS is a close second. Note that the clustered technologies outperform Capillaroscopy, Photo-acoustic and Doppler ultrasound in both plots.

There is now sufficient information to reach a decision.

- If it is desired to reach a solution in the minimal possible time and with minimal risk, and we can accept the best available technique even if the data are quite poor, then off-the-shelf pulse oximeters should be used and the "perfusion index" metric used.

- If we are prepared to develop hardware, but want low risk and to take a route that involves well developed technologies, then it is best to first develop or acquire a continuous wave NIRS system, and then split the development effort between tissue oximetry, venous occlusion, capillary refill and photo-plethysmography, maximizing the probability that a successful device will be produced by reusing the core technology between these techniques.

- Only if we are interested in maximising the potential return on the development effort, and can accept a high risk of the technique proving to be unsuitable for real world use should we pursue thermal monitoring techniques.

Continuous wave NIRS is a fairly clear winner, so it was decided to pursue the NIRS route. First developing the core technology (i.e. CW NIRS), and then experimenting with tissue oximetry, venous occlusion, capillary refill and photo-plethysmography in a effort to devise the most suitable technique amongst them.

7.2 Development of a continuous wave NIRS system

Continuous wave NIRS is available off-the-shelf, but in order to adequately research all the possible techniques, it would be useful to have access to design and firmware documentation in order to for example construct custom made sensors. In the case of capillary refill in particular, simultaneous measurement and or control of applied pressure and data acquisition would be extremely useful. Off-the-shelf NIRS systems are almost always minimally documented with a limited range of end uses, primarily to safeguard intellectual property. The exception usually being University developed systems, such as the system developed at Tufts University described by Seigel et al [214], but then the system is invariably not being manufactured, at least not in its original form. Despite an extensive search no NIRS system with fully documented hardware and an open firmware source was found on the market.

For these reasons it was decided to develop a custom continuous wave system, or at least one with an easily understood flow of data. Near infrared spectroscopy in reflection mode has been carried out in many previous studies using a fibre coupled spectrometer in conjunction with a halogen light source and “pencil” reflectance probe. Off-the-shelf fibre coupled spectrometers are an attractive route to take, but suffer from several disadvantages;

- The path length is very short - less than 1mm. This lowers the sensitivity, and increases the influence of the upper skin layers.

- Spectrometers (e.g. those available from Thor labs and Ocean optics) have a reasonable quantum efficiency, but the very small input aperture adversely effects signal to noise.
- If the path length is increased with a custom made fibre probe, in an effort to increase sensitivity, there is a much lower received light intensity.

It is possible to increase the input intensity, and high intensity fibre coupled light sources are available off-the-shelf. However this risks injury to the subject due to the very high optical flux around the input fibres causing tissue heating. The spectrometer exposure time cannot be significantly increased, or temporal resolution will be lost, meaning that use in PPG will not be feasible. Binning of adjacent spectrometer bins is one possible solution to these problems, giving some improvement to signal to noise, but this technique was found to be insufficient beyond around 2mm separation in simple laboratory tests.

In order to study the PPG signal, or the response of tissue absorption to changing pressure over a time-scale of tens of milliseconds as part of a capillary refill experiment, these problems become insurmountable. A different technique is clearly needed if we are to conduct NIRS over ten millisecond time-scales and source detector separations of several mm.

Use of a monochromator was the route taken by Crowe et al in their PPG studies. By using a 70W high intensity light source followed by a monochromator with 4nm line width, it was possible to obtain a reasonable coupling of several tens of mW into a fibre bundle. But this incident power is still much lower than with an off-the-shelf fibre optic "pencil" reflectance sensor, solving the problem of local tissue heating. More importantly, this method allows a wide area photodiode to be employed as a detector, as only one wavelength will be incident at any one time. This leads to an improvement in signal to noise of several orders of magnitude (around 3), enough to accurately measure the PPG signal. Unfortunately, as evident from the poor quality data presented in the Crowe et al paper [191], a monochromator can only record attenuation at a single wavelength at any point in time, leading to serious artefact problems.

Thus it can be seen that we need to move to a more complex CW-NIRS technology than either spectrometers or monochromators. Almost all the commercial systems use clusters of light emitting diodes (LEDs), with different LEDs emitting at discrete wavelengths. This leads to a potential problem of isolating the detected signals from individual LEDs at the photo-detector(s). A trivial technique would be to use multiple photodiodes with filters, but this is not a very feasible solution - it is much more practical to be able to

use a single photodiode, leaving two possible options:

1. Time division multiplexing - pulse the LEDs in sequence
2. Frequency division multiplexing - modulate the LEDs at separate carrier frequencies.

This leads to a trade off of time division versus frequency division, and a consideration of three factors:

1. Noise (e.g. thermal noise)
2. Interference (e.g. light sources, compact fluorescent lights etc)
3. Impact of artefacts (this is especially relevant given that capillary refill studies involve inducing relatively sudden changes in tissue properties)

Considering these points:

1. Thermal noise has a flat spectrum, but amplifier noise is typically close to 1/f. If it is possible to operate at a high carrier frequency (several kHz) with frequency division, it is possible to avoid a lot of the noise. With time division there is a greater tendency to pick up a lot of the low frequency noise.
   We also need a wide bandwidth front-end to ensure a short settling time with time division. Averaging techniques can reduce the noise introduced by this, but this is a problem that complicates the analogue design. With frequency division we only need good amplifier performance over a narrow bandwidth.
   Frequency division is the optimal solution to this problem.
2. Interference has both a 1/f component from shadows and ambient light changes, and also optical noise sources such as LCD monitors and compact fluorescent lights, which typically have peaks in <5kHz range. Frequency division wins again here as the carrier frequencies can be over 10kHz, above the majority of the optical noise.
3. We require the output data rate to be sufficiently fast that aliasing is minimized. Considering the typical time scale of the pulsatile signal and of capillary refills, this means that the output data rate of the demultiplexing DSP needs to be of order 100Hz or thereabouts.
   However, given the transient nature of the signals, it seems likely that some aliasing will still be experienced at a high output data rate.

Time division will lead to mapping of artefacts more strongly to some channels than others, whereas frequency division will see broadening of the modulation spectrum peaks, giving cross-talk between channels.
Cross-talk is a more preferable effect, so frequency division wins here.

Frequency division was chosen for further development as it clearly gives the best performance.

The carrier frequencies need to be above the corner frequency of the amplifier, and preferably above 5kHz, as already discussed. Many articles have been published on the subject of low noise photodiode amplification, and the most common solution is to use a JFET voltage follower combined with a low voltage noise op-amp in a trans-impedance configuration. This configuration allows use of a low voltage noise but high input current noise JFET op-amp (e.g. Texas Instruments LMH6624). Unfortunately a JFET voltage follower combined with an op-amp such as the LMH6624 leads to a relatively high current consumption and requires an operating voltage of 5V or more, requirements that are not optimal for a portable battery powered device. A parametric search using manufacturers product selection tools shows no low quiescent current, low voltage op-amp with a voltage noise in the $< 5nV/\sqrt{Hz}$ region.

An alternative strategy is to use the junction capacitance of the photodiode to accumulate charge (in a similar fashion to a single pixel of a CMOS image sensor), and either measure the voltage using an op-amp or use the photodiode as part of an op-amp feedback network (figure 7.3). This eliminates the resistor found in the trans-impedance configuration, and thus a source of thermal noise. However, more complex signal processing is required, as the output will be the integrated photocurrent, and there may be some non-linearity. The op-amp input voltage and current noise will still contribute towards the output signal, even though there is now no thermal noise.

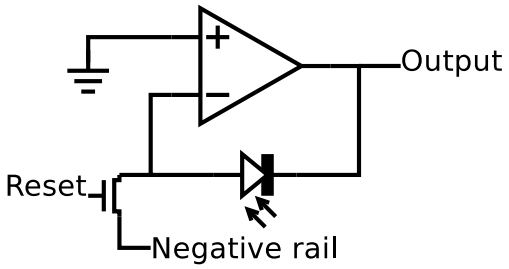


Figure 7.3: A simplified schematic of one potential charge accumulation design.

However, if the amplifier is limited by photocurrent shot noise, then the effect of reductions in amplifier voltage and input current noise will be minimal. In order to make a decision on necessary noise reduction steps, the likely photocurrent will need to be calculated. By assuming credible device parameters and using the diffusion approximation (e.g. Patterson et al [215]), it is possible to calculate the expected photocurrent.

Figure 2 from the publication by Patterson et al [215] plots calculated attenuations for a range of typical tissue properties. It can be seen that a reflectance of around $10^{-5}mm^{-2}$ would appear to be an appropriate worst case reflectance for a 1cm source to detector separation. Assuming a large area photodiode of $10mm^2$ area, a 20mA mean LED current, and a total quantum efficiency (source $\times$ detector) of 50%, then the worst case photocurrent will be in the order of 1μA. This photocurrent was set as the operating point for the front end noise analysis. For reasonable quality photodiodes such as the Vishay BPW34S the dark current is so low, at less than $1nA$, that it can be considered as negligible in the noise analysis.

The photocurrent shot noise can easily be calculated using the electron charge and the standard result for noise in a Poisson process:

$$\sigma_I = \sqrt{2qI\Delta f}. \tag{7.5}$$

Where q is the charge of an electron ($1.6\times 10^{-19}(2.s.f)$) and Δf is the bandwidth. Substituting in our 1μA photocurrent:

$$\sigma_I = 5.7 \times 10^{-13}\sqrt{\Delta f} = 570fA/\sqrt{(}Hz). \tag{7.6}$$

This is the fundamental low limit of the achievable noise; we can lower the front end noise below this level, but the shot noise will always be present due to the stochastic nature of the current flow from the photodiode.

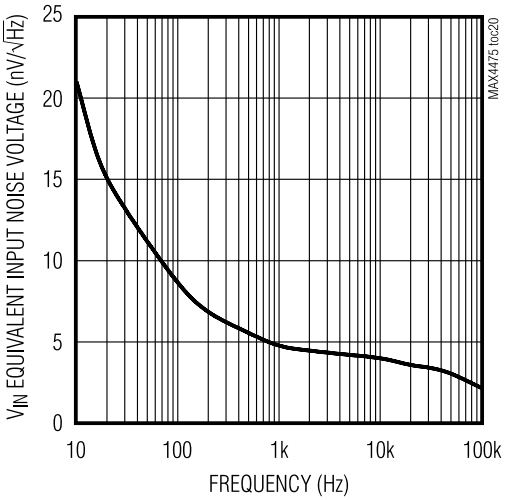


Figure 7.4: Voltage noise density of MAX4477 op-amp (©Maxim Inc.).

It is notable that all available op-amps with input voltage noise of $2\text{nV}/\sqrt{Hz}$ or less have input current noise in the $pA/\sqrt{Hz}$ region, and so do not appear to be ideal solutions unless combined with a discrete JFET input stage to lower the input current noise. In order to model the effect of input voltage noise more accurately, the effective parallel capacitance of the photodiode and cable must be estimated. Using published capacitance for the BPW43S photodiode and assuming 1m of Mogami 2754 miniature shielded microphone cable, the parallel capacitance was estimated as 200pF. Input impedance Z can then be found:

$$|Z| = \frac{1}{2\pi fc} = 7.96 \times 10^8/f. \tag{7.7}$$

Leading to an effective photodiode noise current of:

$$\sigma_{nv} = nf\sqrt{\Delta f}/(7.96 \times 10^8), \tag{7.8}$$

where n is the op-amp voltage noise per $\sqrt{Hz}$.

Assuming an operating frequency of 10kHz, then:

$$\sigma_{nv} = n\sqrt{\Delta f}/(7.96 \times 10^4). \tag{7.9}$$

If this is set to $400fA/\sqrt{Hz}$, slightly under the shot noise, then solving for n gives $n = 32nV/\sqrt{Hz}$. This is a relatively high voltage noise for an op-amp, and shows that use of a BiCMOS amplifier is feasible. Another search using parametric product tools found the Maxim MAX4477 to be an ideal op-amp, with a voltage noise density of $4\text{nV}/\sqrt{Hz}$ and current noise density of $0.5\text{fA}/\sqrt{Hz}$ at 10kHz (figure 7.4). Operation is possible down to a 2.7V supply voltage and the quiescent current is only 2mA, resulting in extremely practical battery powered devices. The gain-bandwidth product is also high enough at 10MHz for effective signal amplification at frequencies in the region of 10kHz.

With a suitable operational amplifier identified, it is still not clear if a trans-impedance design is suitable. This question can be addressed with a consideration of the thermal noise current density from a feedback resistor. Assuming a 3.3V total supply voltage split to give a positive and negative 1.65V rail, a 1MΩ feedback resistor would appear to be the largest suitable size to avoid saturating the amplifier:

$$\sigma_R = \frac{\sqrt{4K_BTR\Delta f}}{R} = \sqrt{4K_BT\Delta f/R}. \tag{7.10}$$

Where K_B is the Boltzmann constant ($1.38 \times 10^{-23} JK^{-1}$) and T is temperature (approx 300K). Substituting in these values:

$$\sigma_R = 1.29 \times 10^{-10}\sqrt{\Delta f/R}. \tag{7.11}$$

Solving this equation for a current noise density of $400fA/\sqrt{Hz}$ gives $R = 103k\Omega$. This resistor size is below the 1MΩ limit and so the shot noise limit can be reached with a simple trans-impedance amplifier configuration.

It is also necessary to consider contact noise from the feedback resistor. As shown by Amaboldi et al [216], contact noise in carbon composition resistors drops as 1/f or greater with increasing frequency. At 10kHz frequency and low voltages, the contact noise is well below the thermal noise for the carbon compound resistors (the noisiest resistor type), so it can be concluded that it is not necessary to use a metal film resistor to reduce noise.

Although all of the noise sources are below the shot noise limit, there are large differences in their magnitudes. The op-amp input current noise density is three orders of magnitude below the shot noise, the current equivalent noise from the op-amp input voltage noise is approximately an order of magnitude below the shot noise, but the thermal current noise

is only approximately three times smaller than the shot noise. There might therefore be a modest performance improvement from reduction of the thermal current noise from the trans-impedance resistor.

Three steps that can be taken to give a modest reduction in thermal noise are: 1) use a single rail or unevenly split rails to give more headroom for the op-amp, 2) use a high feedback capacitance so that at the operating frequency, the complex feedback impedance Z is primarily imaginary. This results in the AC component of the photo-current being suppressed, so the op-amp headroom only needs to be sufficient for the DC component, 3) increase the supply voltage to give more headroom and allow the resistor size to be increased, this is less appealing as higher voltages are less practical for simple battery powered apparatus.

A detailed analysis of the MAX4477 datasheet shows that the first option will not result in stable operation with a 3.3V supply voltage, but the second option is feasible with the second op-amp in the dual amplifier SOIC-8 MAX4477 being used as an active bandpass filter/amplifier to bring the AC voltage up to the full ADC range. This strategy has the additional advantage that there is no DC offset at the ADC, allowing a lower resolution ADC to be used than would be the case with a DC coupled trans-impedance amplifier.

Power supply rejection ratio (PSRR) has not been considered in this analysis. The MAX4477 PSRR is 50dB at 10kHz, meaning that power supply noise would have to be of order $1\mu V/\sqrt{Hz}$ or greater to be significant. As many low noise voltage regulators have voltage noise of less than $100nV/\sqrt{Hz}$ at 10kHz, it would appear that power supply noise will be negligible if an appropriate power supply design is used.

7.2.1 Front end electronic design - dual op-amp system

The design concept previously outlined of a trans-impedance amplifier stage followed by an AC coupled (via a 100nF capacitor) active band pass gain stage was developed for the photodiode front end amplifier. A two stage prototyping procedure was used to develop the design, before a "final" version of the design was constructed on FR4 based strip-board for later use in the DSP and firmware development.

The first prototyping stage was simulation based, with the LTSpice package (a freeware SPICE package supported by Linear Technology Inc.) being used to simulate the circuit and fine tune component values to obtain optimal performance. The chosen design was then prototyped on solderless breadboard. An LED driven by a square wave from a signal generator was used to illuminate the photodiode, and the output voltage monitored with an oscilloscope to compare the performance to the SPICE simulation.

In order to maximise the out of band rejection prior to the filtered AC signal entering the ADC, an Infinite-Gain Multiple Feedback (IGMF) active filter design was chosen for the second stage of the amplifier. This allowed a high Q with a single op-amp, so the entire front end could be based around a single MAX4477 dual amplifier IC.

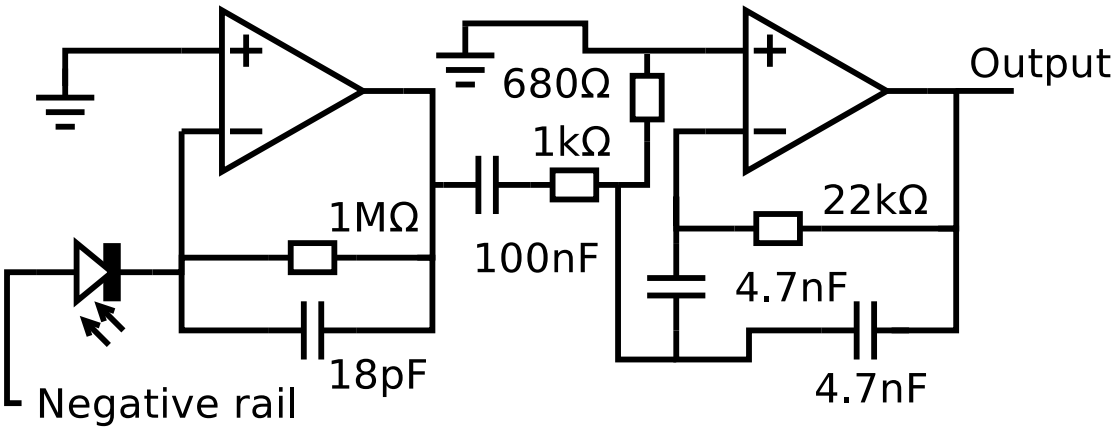


Figure 7.5: The final design for the photodiode front end amplifier, adopted after SPICE and breadboard prototyping.

Figure 7.5 is a schematic of the final design. The rail splitter is not shown (it consisted of a potential divider using two 10kΩ resistors, and a 4.7μF and 100nF decoupling capacitor). It was found that a 2.2μF capacitor was sufficient for decoupling the virtual ground in SPICE, but this resulted in a complex oscillation with an unstable period in the millisecond range when the amplifier was breadboarded. This was found to arise due to current being injected into the virtual ground via the active bandpass filter, causing biasing of the midpoint voltage. A larger 4.7μF capacitor was found to result in reliable performance with the breadboarded prototype.

Although it is theoretically possible to create a low gain, high Q bandpass filter using the IGMF filter, in practice this combination results in unfeasible component values, and SPICE simulation indicated asymptotically increasing susceptibility to both parasitic current leaks (e.g. op-amp input bias current) and component tolerances. For this reason the overall amplifier gain was increased, meaning that shot noise will be more significant than the previously calculated value. The peak gain frequency was also increased from 10kHz to 11.6kHz, making component selection easier. SPICE simulation indicated that a $\pm$1% change in resistor values and a $\pm$10% change in the value of the 4.7nF capacitors resulted in a worst case change in gain at 11.6kHz of 15%, which was judged to be acceptable. Assuming operation in "office" conditions (i.e. 10 to 35°C), this means that extremely common and low cost 5% tolerance X5R capacitors can be used. Standard carbon composition resistors have 1% or better tolerance, so component selection is not demanding.

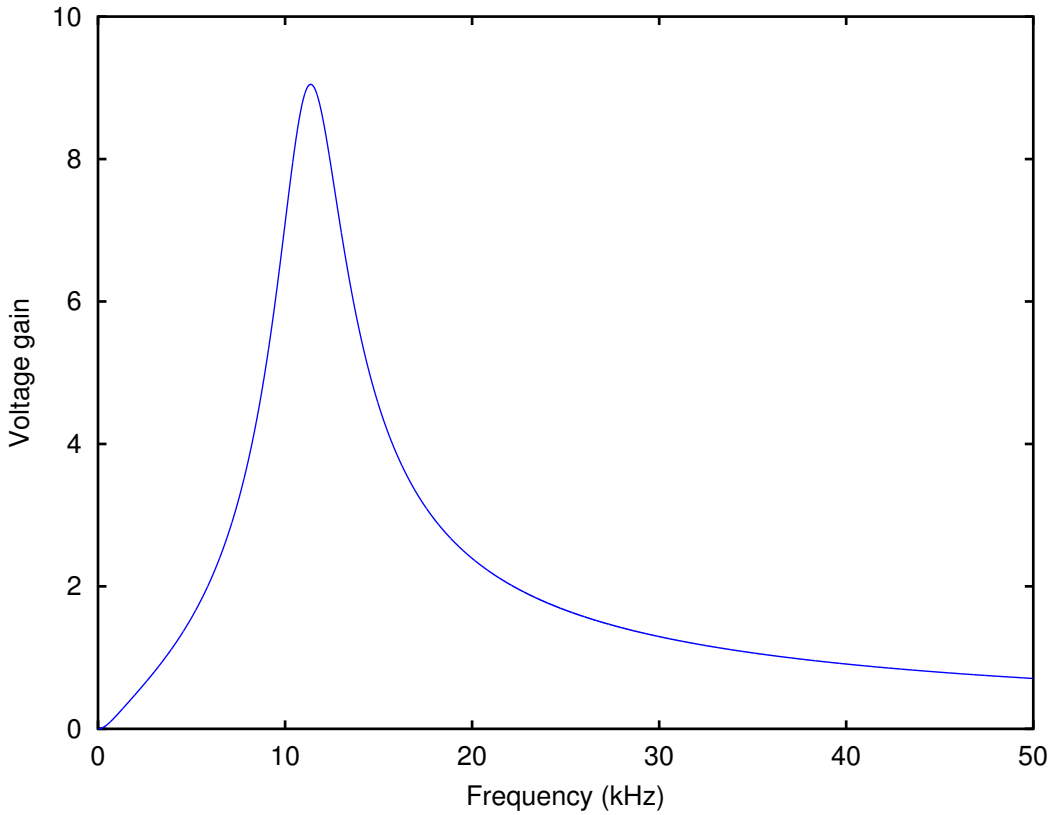


Figure 7.6: Voltage gain of the second stage filter - simulation using SPICE package.

An 18pF capacitor was placed in parallel with the 1MΩ resistor in the first stage trans-impedance amplifier. This results in a trans-impedance that is primarily capacitive at the 11.6kHz operating frequency, decreasing the AC gain and thus allowing a larger photocurrent without the thermal current noise of a small value resistor (a more detailed explanation of this effect was given in the previous section). The trans-impedance capacitor also acts as a first order low pass filter at frequencies significantly greater than the 8.8kHz transition frequency of the trans-impedance stage, leading to significantly more attenuation at higher frequencies than plotted in figure 7.6, and thus a further reduction in high frequency noise injected into the ADC. This is an important design improvement, as noise at frequencies above the Nyquist could potentially alias down, causing interference at the modulation frequency.

By making some assumptions concerning the parasitic capacitance of the photodiode cable, the junction capacitance of the photodiode, and the voltage and current noise of the op-amp (with values taken from MAX4477 datasheet), the effective photocurrent noise density of the final amplifier design shown in figure 7.5 can be calculated. The current noise equations for the 1MΩ resistor thermal current noise and the current noise resulting from the op-amp voltage noise acting across the parasitic and junction

capacitances have been calculated previously, and the MAX4477 input current noise is small enough to be considered insignificant. Assuming a 3.3V peak to peak output voltage into the ADC, the photocurrent can then be calculated. At 11.5kHz, the trans-impedance will be:

$$1/Z = 1/10^6 + 1/\left[\frac{-j}{2\pi \times 18 \times 11.6 \times 10^{-9}}\right] = 10^{-6} + (1.3 \times 10^{-6})j, \tag{7.12}$$

$$|Z| = 6.1 \times 10^5\,(2s.f.). \tag{7.13}$$

Assuming a voltage gain of 9 times in the second stage amplifier, the AC voltage at the trans-impedance amplifier output will be 366mV peak to peak. Assuming a sinusoidal photocurrent and dividing by the trans-impedance Z, the peak to peak current will be 600nA (2s.f). With negligible ambient light, the mean photocurrent will be 300nA, and the shot noise density will be approximately $\sigma_I = \sqrt{2qI\Delta f} = 310fA/\sqrt{Hz}$. Note that this analysis is somewhat simplified as the shot noise will rise and fall with the square root of the sinusoidal photocurrent. This effect can be accounted for by scaling by a factor of:

$$\frac{\int_0^\pi \sqrt{1 + cos(\theta)}\,\mathrm{d}\theta}{\pi} = 2\sqrt{(2)}/\pi = 0.90\,(2s.f.). \tag{7.14}$$

Thus the shot noise is $280fA/\sqrt{Hz}\,(2s.f.)$.

Substituting a value of 1MΩ into equation 7.11, $\sigma_R = 130fA/\sqrt{Hz}\,(2s.f.)$. Likewise, substitution into equation 7.8 leads to $\sigma_{nv} = f \times 4.2 \times 10^{-9}\sqrt{\Delta f}/(7.96 \times 10^8) = 61fA/\sqrt{Hz}$. Putting all these figures together, an estimate of total photocurrent noise density can be arrived at:

$$\sigma_{Tot} = \sqrt{\sigma_R^2 + \sigma_{nv}^2 + \sigma_{shot}^2} = 315fA/\sqrt{Hz}. \tag{7.15}$$

It can be seen that shot noise accounts to 89% of the total noise, so we are effectively shot noise limited.

With an estimate of the photocurrent noise density, it is now possible to calculate the signal to noise ratio. The peak to peak photocurrent was previously calculated as 600nA, giving an AC RMS photocurrent of 210nA (2s.f.) thus the signal to noise ratio will be:

$$SNR = \frac{210 \times 10^{-9}}{315 \times 10^{-15}\sqrt{\Delta f}} = 6.7 \times 10^5/\sqrt{\Delta f}. \tag{7.16}$$

This can also be expressed as an effective number of bits (ENOB) per intensity sample, a useful form for aiding the choice of ADC and DSP:

$$ENOB = \frac{ln(6.7 \times 10^5/\sqrt{\Delta f})}{ln(2)} = 19.3 - 0.72ln(\Delta f). \tag{7.17}$$

If a noise bandwidth of 50Hz is assumed, sufficient for both the pulsatile blood volume fluctuations and transient effects associated with capillary refill experiments to be observed, then the ENOB will be $19.3 - 0.72ln(50) = 16Bits\,(2s.f.)$. Considering that typical near infrared pulsatile modulation depths in reflection mode are in the order of 0.1%, it can be seen that around 6 useful bits of data will be present when observing the PPG signal, more than enough to be useful.

7.2.2 Front end electronic design - single op-amp system

The system outlined in the previous section performs well since it is shot noise limited, meaning that increasing the light source intensity would be the most effective method to increase SNR. This comes at the expense of LED drive current, and the current is already in the order of 20mA, which is more than the amplifier and an attached low power digital signal processor would draw. Thus it can be seen that in the case of a portable device, the amplifier design places the design optimisation problem in the SNR versus battery life territory, a fundamental trade off caused by the quantisation of photocurrent electron charge, that cannot be avoided through any more amplifier optimisation.

However, there are still two ways in which the amplifier design could be seen to be sub optimal:

1. If the light intensity is too high, the first stage trans-impedance amplifier will saturate. There is no easy way to detect a saturating front end unless an additional ADC is used to measure the output of the trans-impedance stage or a comparator is placed there and the output logic level used as an indicator. The MAX4477 running at 3.3V requires an evenly split rail, so there is only 1.65V of headroom before the output will saturate.

2. The component count is relatively high, and the dual amplifier IC is quite large, taking up valuable board space in a portable device.

For these reasons a single op-amp design was developed using the same procedure of SPICE followed by breadboarding. In order to exclude as much noise from the ADC as possible, it would be optimal for the single op-amp design to incorporate a high Q filter.

Additionally, if a trans-impedance design is used, the magnitude of the impedance at the optical modulation frequency must be very large in order for the sensitivity to be sufficiently high. These requirements suggest that a resonant feedback network would be suitable, but a careful consideration of the component values required for a parallel inductor and capacitor to be resonant at 10kHz shows that this is not practical for reasonably sized surface mount components.

However, it is possible to simulate high value inductors using a gyrator circuit. This opens up the possibility of creating a resonant feedback network if such a gyrator can be combined with a parallel capacitance to form the feedback impedance.

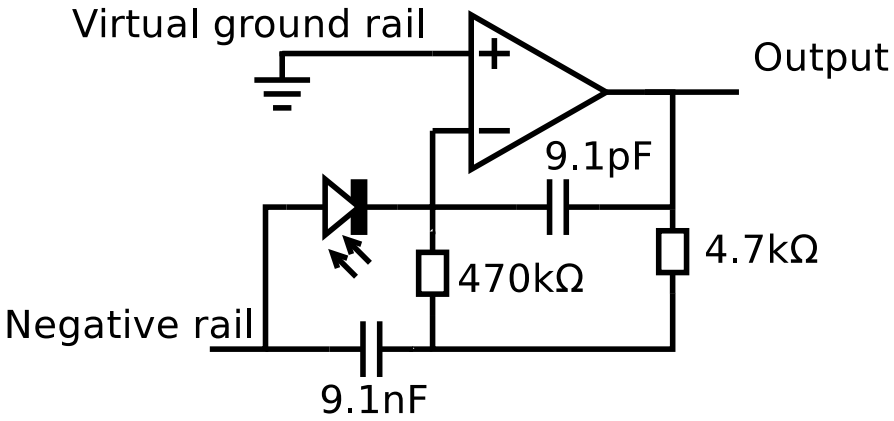


Figure 7.7: Single amplifier photodiode amplifier, designed using optimisation in SPICE package

Although it is theoretically possible to calculate the effective inductance of a simulated inductance in a gyrator circuit, in practice it was found that the behaviour became considerably more complex when a parallel capacitor was incorporated into the design. For this reason SPICE was used to iteratively improve the design until resonance was at around 11.6kHz with both a high Q and a high magnitude trans-impedance. Figure 7.7 is a schematic of the final design. As with the previous circuit, this was then prototyped, but it was necessary to use FR4 based stripboard due to the high sensitivity to parasitic capacitances in this design. The prototype was found to match the simulation well. An analysis of the effects of component tolerance and parasitic op-amp bias current was conducted in SPICE, and it was found that ±1% tolerance was acceptable for all the resistors, and ±10% for the 9.1nF capacitor, meaning a ±5% X5R ceramic capacitor could be used. However, the circuit was found to be very sensitive to variation in the size of the 9.1pF capacitor, with the resonant frequency moving sufficiently for a ∼40% shift in sensitivity to occur with a 10% change in capacitance, meaning the X5R type is likely to be unsuitable and a higher quality ceramic capacitor is required, for example NP0 with a tolerance of ±2% or better. It can be seen that the thermal current noise

will originate from the 470kΩ resistor, resulting in a slightly higher thermal current noise than in the case of the previous design. It can also be seen that the virtual ground serves only to set the operating level of the op-amp, so the decoupling is not nearly as challenging, and the virtual ground is not required to source or sink significant current, meaning that a high impedance DAC with decoupling sufficient to remove only the thermal noise (a 100nF or less capacitor would suffice) would be a sufficient source for the rail.

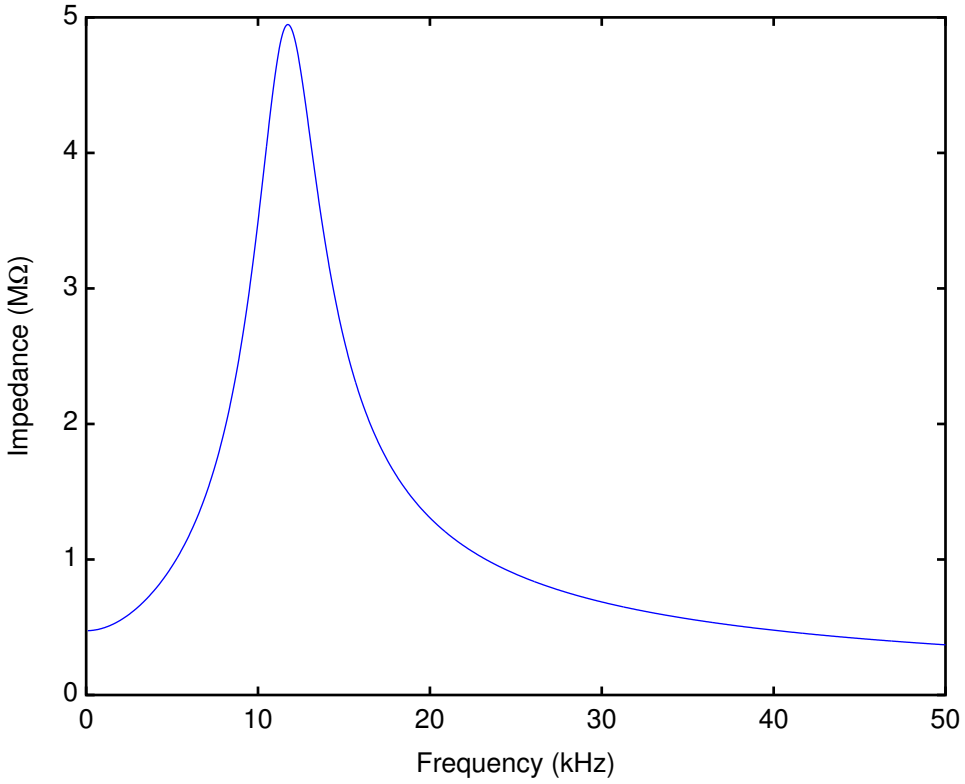


Figure 7.8: Trans-impedance of the single amplifier photodiode amplifier

Figure 7.8 shows the effective trans-impedance. The resonant peak is approximately 5MΩ at a frequency of 11.8kHz. Comparing to the previous design, it can be seen that the peak gain is approximately 10% lower (from equation 7.13, applying a gain of 9 times gives 5.5MΩ). The rejection at high frequencies is also moderately lower, and the current gain at 20kHz is three times higher for the single op-amp design than the dual, and once the effect of the 18pF trans-impedance capacitor is accounted for there is a first order cut off at even higher frequencies, compared to a second order cut off from the previous design. Another important difference is the DC coupling, meaning that saturation of the front end due to excessive photocurrent can be detected by monitoring the DC component of the ADC signal; i.e. simply looking for a saturated ADC.

Due to the DC response, it would be optimal to run this circuit with a virtual ground that is lower than the midpoint between 0 and 3.3v. Unfortunately, as previously discussed, the MAX4477 op-amp is not stable in this configuration, so another opamp was sought using manufacturer parametric product tools. The MCP6021 from MicroChip Inc. would appear to be the most suitable replacement. This op-amp has the additional advantage that it is available in a small 5 pin SOT23 package. However, the voltage noise is larger than that of the MAX4477, meaning that it is necessary to repeat some of the noise analysis conducted previously for this new design.

If the DC trans-impedance is 470kΩ and the device is operating at resonance, so that the AC trans-impedance is 5MΩ, then it is possible to calculate the virtual ground voltage that would give the maximum possible modulation depth without clipping the top rail. However, in real world use there will also be some ambient light passing through the tissue and striking the photodiode. In a series of volunteer tests under typical office lighting conditions, this illumination was found to give an absolute maximum photocurrent of 500nA, giving a DC offset of 235mV. Thus we have a system of equations that can be solved to find the optimal virtual ground voltage:

$$3.3 = (4.7 \times 10^5 \times I_{photo}) + 0.235 + (5 \times 10^6 \times I_{photo}) + V_{virtual\_ground}, \tag{7.18}$$

for the peak voltage, which can simplify to:

$$3.065 = (5.47 \times 10^6 \times I_{photo}) + V_{virtual\_ground}. \tag{7.19}$$

The lowest voltage must remain above the ground rail, so we have:

$$0 = (4.7 \times 10^5 \times I_{photo}) + 0.235 - (5 \times 10^6 \times I_{photo}) + V_{virtual\_ground}, \tag{7.20}$$

which can be simplified to:

$$4.53 \times 10^5 \times I_{photo} = V_{virtual\_ground} + 0.235. \tag{7.21}$$

So there is a system of two simultaneous equations, which can be solved to give the maximum photocurrent. Solving equations 7.19 and 7.21 using a 2×2 matrix inversion method gives $V_{virtual_ground} = 1.26$V. It is now possible to solve for the photocurrent by dividing by the 5MΩ trans-impedance, giving $I_{photo} = 250nA\,(2s.f.)$, or a peak to peak photocurrent of 500nA, slightly smaller than with the previous amplifier design.

The photocurrent shot noise density can be found using equation 7.5, and applying the scaling factor (equation 7.14) to account for the sinusoidal photocurrent, gives $\sigma_{shot} = 250fA/\sqrt{Hz}\,(2s.f.)$. As in the previous design analysis, the trans-impedance thermal current noise can be calculated from the trans-impedance (in this case the 470kΩ

resistor) and the effect of voltage noise from the op-amp specification (8.5nV/$\sqrt{Hz}$ at 11.8kHz). Applying equation 7.10, the thermal current noise can be found as $\sigma_R = 190fA/\sqrt{Hz}\,(2s.f.)$, and applying equation 7.9 to the 5.0MΩ trans-impedance, the effect of op-amp voltage noise can be calculated as $\sigma_{nv} = 126fA/\sqrt{Hz}$, leading to:

$$\sigma_{Tot} = \sqrt{\sigma_R^2 + \sigma_{nv}^2 + \sigma_{shot}^2} = 338fA/\sqrt{Hz}. \tag{7.22}$$

So the total photocurrent noise density is $340fA/\sqrt{Hz}\,(2s.f.)$. This is significantly higher than the shot noise density, which is 26% lower than the total noise, so the single op-amp amplifier is not quite at the shot noise limit, but very close. The SNR and ENOB figures can now be calculated, in a repeat of the procedure for the previous design. Using an RMS AC photocurrent of 178nA:

$$SNR = \frac{178 \times 10^{-9}}{338 \times 10^{-15}\sqrt{\Delta f}} = 5.3 \times 10^5/\sqrt{\Delta f}, \tag{7.23}$$

$$ENOB = \frac{ln(5.3 \times 10^5/\sqrt{\Delta f})}{ln(2)} = 18.9 - 0.72ln(\Delta f). \tag{7.24}$$

Comparing to equation 7.17, it can be seen that the single op-amp design sacrifices less than half a bit of effective data per sample. The simpler, more stable and lower quiescent current single opamp circuit would appear to be the best choice unless absolute lowest noise is required. The single op-amp design also avoids the potential for undetected front end saturation to occur.

7.2.3 DSP design

With the frontend design thoroughly explored, the next stage in the design of the NIRS system was the ADC and DSP. So called "Digital Signal Controllers" (DSC) are now widespread. These devices marry a set of conventional micro-controller peripherals, such as pulse width modulators (PWM), analogue to digital converters (ADC), serial peripheral interface busses (SPI) ect. with a high speed core featuring DSP extensions. The ARM Cortex™ M4F is a popular core in industry, using the common Thumb2 instruction set with DSP instructions such as single cycle 32 bit MAC. The "industry standard" ARM architecture is supported by many toolchains, e.g. the GNU GCC toolchain. The M4F is also the only M (i.e. "micro-controller") suffix ARM Cortex™ core to feature an FPU. At the time that the system was developed (2011), the ST-microelectronics STM32F4 range was one of the few available Cortex M4F micro-controllers. The STM32 range also feature a high performance 12 bit ADC. This allows a sampling rate of up to

1Msps, and unlike the ADC on almost all similar micro-controllers, the effective number of bits does not fall significantly as the sample rate is increased up to the maximum. In addition the entire STM32 range features flexible high performance 16 and 32 bit timers with features such as gating and set/reset on routable internal signals. This may allow the LED modulation to be generated using the on chip timers. For these reasons the STM32F4 micro-controller was chosen for use in the prototype development.

The "F4 Discovery" evaluation board is available from ST-microelectronics, and features a 100pin STM32F4 with all pins broken out to 0.1 inch headers, a USB interface, voltage regulator, indicator LEDs and a built in SWD/JTAG debug interface. Due to the ready availability and easy of use, this development board was chosen for use in the initial firmware development. A daughterboard consisting of a MAX4477 based dual op-amp front end (for maximum signal to noise ratio), an analogue rail filter (consisting of a multi-stage pi filter to remove high frequency noise from the 3.3V digital rail tapped from the discovery board), and a series of MOSFET transistors to drive the LEDs was constructed on FR4 based stripboard and fitted with 0.1 inch sockets to sit below the F4 Discovery board. Three millimeter holes were drilled to allow the electronics to be bolted to experimental apparatus with M3 bolts.

An RTOS (Real Time Operating System) was used in order to speed development time and make the prototype code more flexible. The ChibiOS RTOS was selected for this purpose, as it is an open source, GPL licensed RTOS with a built in hardware abstraction layer which supports all the peripherals on the STM32 range of processors. A simple dual threaded application was written to run on-top of ChibiOS, with a background loop passing data to a PC via the USB interface, and a DSP thread demodulating and down-sampling the ADC data, and passing raw intensities to the background loop via a mailbox FIFO. The DMA was used in dual interleaved mode to write ADC data directly to RAM, and trigger an interrupt request on DMA half transfer complete. The DMA interrupt service routine then awoke the DSP thread, causing it to process the recently completed DMA half transfer.

The comprehensive USB drivers built into ChibiOS were used to put the STM32 USB-OTG port into device mode and create a CDC-ACM class interface. This was interrupt driven and accessed from the background thread loop via buffers. The CDC-ACM class can be mounted as a device on Linux based PCs, typically at /dev/ttyACM<x> on the filesystem, where <x> is a number starting from 0 and assigned when the device is mounted. Thus if two devices are connected simultaneously, they will be at ttyACM0 and ttyACM1. The CDC-ACM device appears as a transparent byte stream, allowing bytes to be sent over the USB link to the background thread without the requirement

to write custom USB drivers or handle complex USB packetization. A simple terminal interface was written for the micro-controller using sscanf and printf to handle initialisation settings and intensity data output over the USB.

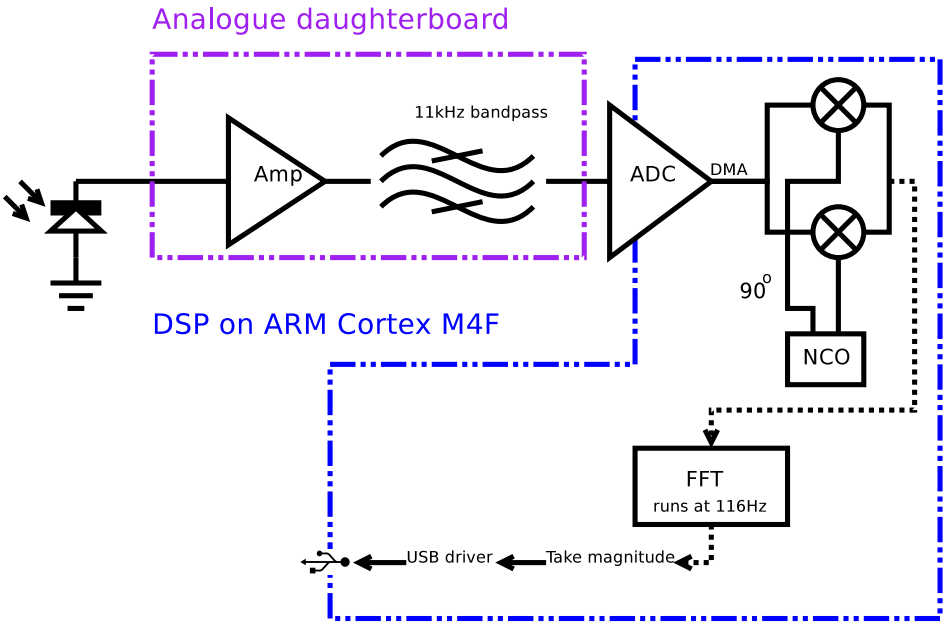


Figure 7.9: The prototype NIRS system, showing the front end and DSP

Figure 7.9 shows a schematic overview of the prototype system. In the DSP thread loop the ADC data are first multiplied by a complex local oscillator generated from a lookup table, giving complex data with a zero intermediate frequency. Due to its reduced number of operations (fewer MAC operations than DFT), a fast Fourier transform was then used to separate the individual carriers. The intermediate frequency data were downsampled before the FFT in order to further reduce the number of CPU cycles required per ADC sample. The final output data rate is $116Hz$, a frequency that was chosen to give sufficient bandwidth for dynamic effects such as the capillary refill and pulsatile PPG signal to be observed with ease. In order for the FFT to separate the individual carriers orthogonally (i.e. with no crosstalk), they must be offset from one another in frequency by integer multiples of the $116Hz$ FFT output data rate. Thus it can be seen that a method of producing multiple carrier with $116Hz$ frequency offsets is required.

There are pulse width modulation outputs on ten of the STM32F4 timers, and it would be appealing to use these to modulate LEDs. As pulse width modulation uses a comparator to set an output pin level based on the current value of a timer, it can be seen that the PWM frequency will be equal to the timer overflow rate. This is set by a timer reload value, so each time a counter reaches the reload value it is reset to zero. On the STM32 all peripherals are driven from a central phase locked loop, so are integer

fractions of the system clock, which in the case of the prototype was set the to maximum $168MHz$ rate to avoid having to optimise code efficiency to any significant extent. It is thereby apparent that the PWM is only capable of generating waveforms with periods of consistent integer offsets, and as frequency is defined as the reciprocal of the period, waveforms with consistently offset frequencies would appear to be an impossibility.

The simplest way around this problem would be to calculate the cross talk between modulation channels, and use the inverse of the mixing matrix to correct the raw data and obtain the true channel intensities. However, when this was attempted with the prototype hardware, it was found that the slowly changing phase lag between the timers and ADC sampling led to a complex ripple sequence, caused by a combination of crosstalk between the STM32 peripherals, and the behaviour of the photodiode front end, which allowed some of the higher harmonics of the PWM modulation to enter the ADC. This ripple contained harmonics in the $1Hz$ range, making the intensity data unusable for most purposes.

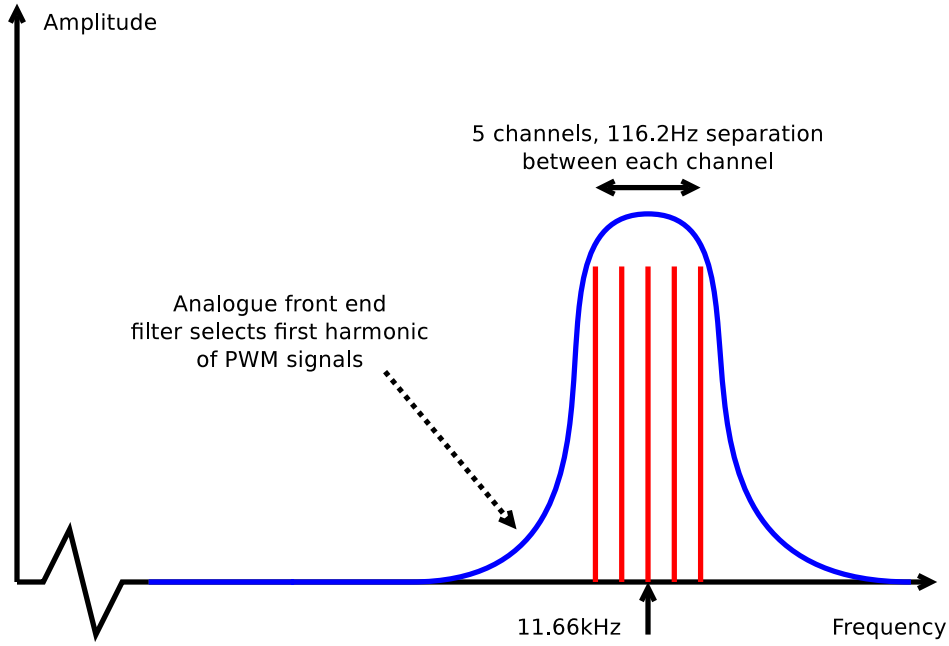


Figure 7.10: The frequency modulation scheme with an idealised front end envelope

There are a number of other ways around this problem. For example, alongside the timers used to generate a PWM, an additional timer could be used to trigger an interrupt request at $116Hz$, with the interrupt service routine resetting the timer counters. This would eliminate the drifting phase, but leave the crosstalk, so the unmixing matrix would still have to be applied. Although this reset technique was found to be usable, and was employed in some of the early firmware development effort, it was found to lead to excess noise in the signal due to the non deterministic nature of interrupt handling

on the STM32. The STM32F4 runs from multiple banks of NOR flash, with a 128bit wide flash bus and instruction and data caches with pre-fetch and branch speculation. This system was found to cause non-deterministic delays on entry to an interrupt service routine placed at addresses in flash memory. In an effort to address problems of non-determinability and high crosstalk, use of the timer gating feature of the STM32 timers was investigated.

The timer gating feature allows a PWM output channel from one timer to be used to selectively enable and disable the clock to a second timer. By use of built in programmable delay lines it is possible to "slave" the clock lines of multiple timers from a PWM output of a "master" timer. As each timer can only have one "master" input and one "slaving" output, it is only ever possible to connect the timers in a chain. This strategy clearly opens up more degrees of freedom from which it may be possible to pick a set of timer reload and gating values that results in the optimal equally spaced frequencies. However, this is a relatively challenging problem to solve, as not only are there timer reload and gating values to be set, but the optimal gating topology must also be found. In order to solve this problem, a simulation of multiple gated timers was written in the GNU-Octave environment, and the topology was optimised with a script to minimized a cost function. This cost was defined as the maximum fraction crosstalk between any pair of channels. If no set of reload and gating values for a given topology could be found that did not result in phase drift between any pair of timers over a period of one FFT iteration, then the topology was discarded. So as to make solutions with no timer phase drift possible, the centre frequency was set to the closest value to $11.8kHz$ that resulted in an integer number of modulation cycles per FFT period. This resulted in the carrier frequency centre of $11.6\dot{6}kHz$.

Figure 7.11 shows the optimised timer topology. It was found that clocking the zero frequency timer first, followed by slaving of the other timers in the frequency increment order +1,-1,+2,-2 resulted in low crosstalk in the 1% or less range whilst maintaining consistent phase for all the timers.

It was judged that a five channel device was sufficient for the planned NIRS research, so although the STM32F4 is capable of driving up to 7 channels in similar timer configurations, only the schematic in figure 7.11 was implemented on the F4 Discovery. With only 5 channels to demodulate, it was found that an FFT is slower than the equivalent DFT, so the DSP code was rewritten to a hardcoded 5 channel DFT rather than using the generic FFT library from the ARM CMSIS™ library.

If the carriers all have equal phase, then the peak photocurrent will simply increase linearly with the number of carriers. Thus it can be seen that in the case of equally

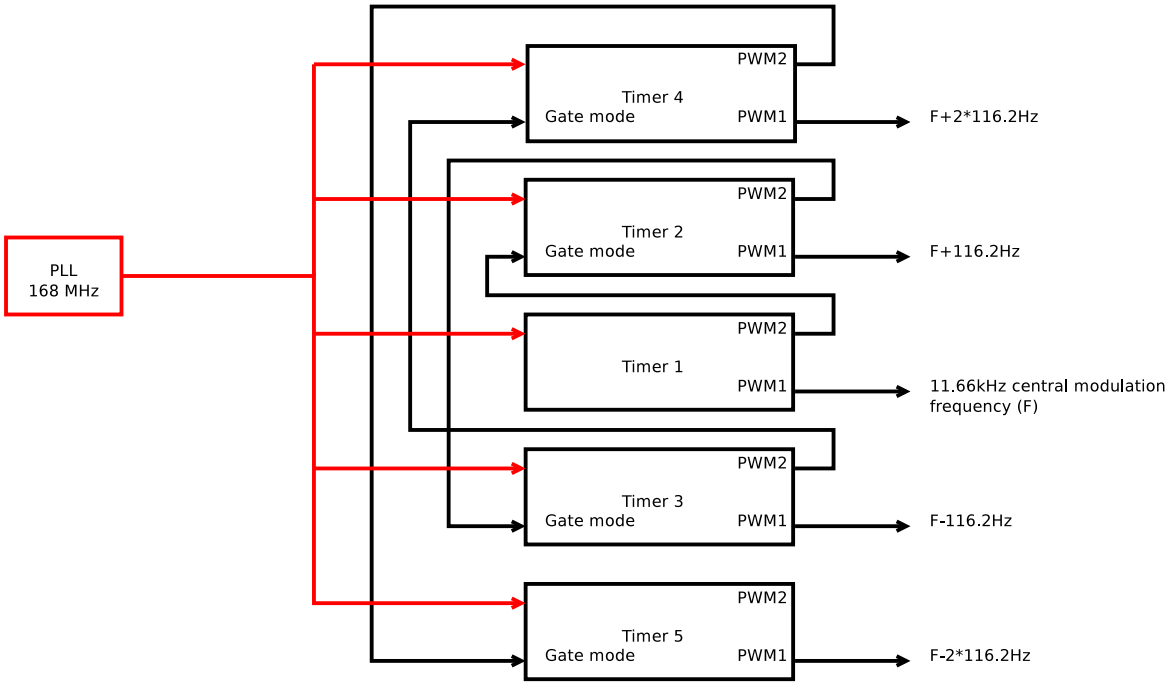


Figure 7.11: The optimised timer and PWM configuration for the frequency modulation scheme

phased carriers, the SNR per channel will decrease in proportion to the number of channels. However, by suitable phase arrangement it should be possible to decrease the peak photocurrent, and therefore increase the mean photocurrent and thus the signal to noise ratio. Minimization of the peak to average power ratio (often abbreviated to PAPR) is a problem that has been studied extensively in the context of orthogonal frequency division multiplexed QAM systems. In the case of communications, the situation is complicated by the variable amplitude and phase, however in a NIRS system the phase will remain constant, so an automatic gain control loop algorithm can be used to adjust the pulse width to each LED channel, resulting in approximately equal photocurrent amplitudes at each modulation frequency, even if optical attenuation and photodetector/LED quantum efficiency differs between the wavelengths. This consistent amplitude and phase makes the NIRS system simpler to analyse, but also means that OFDM solutions derived from QAM are hard to apply.

In order to determine both the optimal phase arrangement, and the relationship between the number of channels and the photocurrent per channel, a numerical model was written in the GNU-Octave environment. The sqp (successive quadratic programming) optimisation function was used to adjust the initial signal phases in order to minimize the peak photocurrent. By appropriate choice of a reference frequency and time, forcing one of the carriers to have a constant zero phase, and a second carrier to have a phase of zero at zero time, it can be seen that the degrees of freedom for optimisation is

equal to the number of carriers minus 2. It was found that the lowest achievable peak photocurrent for a set of n equal intensity orthogonally modulated photocurrents was:

$$I_{peak} = I_{carrier}\sqrt{n+2}. \tag{7.25}$$

This equation is valid for any situation with multiple carriers. In the case of a single carrier, it is clear that $I_{peak} \equiv I_{carrier}$. The signal to noise ratio and thus the ENOB per channel can be calculated by applying equation 7.25 to equation 7.17 (as the MAX4477 front end was used on the F4 Discovery daughterboard). With a bandwidth of 116.2/2Hz for each channel, and with the photocurrent per channel reduced to prevent saturation at the peak photocurrent, this leads to the following equation for the ENOB:

$$ENOB = 19.3 - 0.72ln(58.1) - ln(\sqrt{n+2})/ln(2) = 16.4 - 1.44ln(\sqrt{n+2}). \tag{7.26}$$

In the 5 channel case, this can be solved:

$$ENOB = 16.4 - 1.44ln(\sqrt{7}) = 15.0\,(3s.f.). \tag{7.27}$$

This is enough information to provide useful data for analysis of the pulsatile PPG for example, even if the modulation depth is as low as 0.1%. In practice it was found that the square wave modulation created by the PWM leads to some of the higher harmonics entering the ADC, and increasing the peak voltage. The earlier analysis of the shot noise also assumed a sinusoidal modulation, and PWM modulation leads to a higher shot noise that is dependent on the pulse width - i.e. the source brightness. For these reasons it proved necessary to reduce the illumination intensity by around 20% compared to the theoretical maximum, resulting in the loss of around 0.3 bits per sample, and an ENOB of 14.7 per sample.

7.2.4 NIRS sensor development

Although conventional pulse oximetry sensors are available off-the-shelf, they are limited to two wavelengths. Off-the-shelf near infrared sensors are usually fibre coupled, or use a larger path length than would be practical for investigation of capillary refill. Use of an off-the-shelf sensor would also mean the optical wavelengths would be fixed by the choice of product. For these reasons a custom made reflection mode sensor was constructed with a path length adjustable up to approximately 6mm.

An off-the-shelf flat flexible cable was used to power the LEDs, and a μFL to RP-SMA coaxial cable assembly was constructed for the photodiode signal return. This is also

available as an off-the-shelf part. A sensor arrangement with dual LED sources for each channel was chosen in order to minimize the effect of an uneven or off centre application pressure. This resulted in an outer ring of 20 LEDs, as shown in figure 7.14. Only 10 of the LEDs were designed to be populated in any one sensor, allowing a single PCB design to be used for two different source to detector separations. The outer diameter of the PCB is 18mm, giving a separation of approximately 5.5mm for the outer LEDs and 2.5mm for the inner. LEDs were standard off the self 0603 surface mount devices for the 660nm peak (Kingbright KP-1608SURC) and 950nm peak (Osram SFH4010). The Roithner SMC series were used for all other wavelengths (Roithner Lasertechnic GmbH, Vienna, Austria). Figure 7.12 shows the spectrum of the 5 LEDs. It can be seen that each emitter is relatively broad and there is considerable overlap. If a 7 channel system was to be constructed (as the F4 Discovery board is capable of modulating up to 7 channels with the PWM scheme), then LEDs emitting at around 700nm and 900nm would appear to be suitable for improving the near infrared coverage. Data were recorded with a calibrated Thor Labs CCS200 compact spectrometer (Thor Labs Ltd, Ely UK). Intensity has been converted to arbitrary units, but the area under each curve has not been normalized. The CCS200 uses a Toshiba TCD1304DG linear CCD sensor, which has a dramatically decreasing spectral response past 600nm. Considering that the CCD sensitivity plot (defined by Toshiba as the normalized quantum efficiency) decreases by a factor of 9 between 660nm and 950nm, that photon energy decreases by a fraction of approximately 1/3, and as the PWM pulse width was set to a consistent value for all LEDs whilst the spectrometer exposure was made, it can be seen that the quantum efficiency of the different LEDs is approximately the same (to within $\sim \pm 20\%$).

Figure 7.13 is taken from the VTB8441H datasheet. It can be seen that a quantum efficiency of 0.70 provides a reasonable match to the data between 450nm and 950nm.

In order to meet device standards for bio-medical experiment volunteer safety, the sensor was encapsulated in a transparent silicone rubber with documented bio-compatibility to the relevant parts of the ISO10993 standard for temporary skin contact. A mould for encapsulating the PCB with a 20mm diameter and 6mm deep layer of silicone was machined for this purpose from type 304 stainless steel using electric discharge machining.

The use of transparent silicone encapsulation presented a potential for direct coupling of light from the LEDs to the photodiode. In order to prevent this, the transparent plastic encapsulated Vishay BPW34S photodiode used for early prototyping was changed for a Excelitas VTB8441H, which is encapsulated in a light absorbent black ceramic with an epoxy resin window over the active area of the die face only.

Due to the extremely low surface energy of the two component silicone rubber encap-

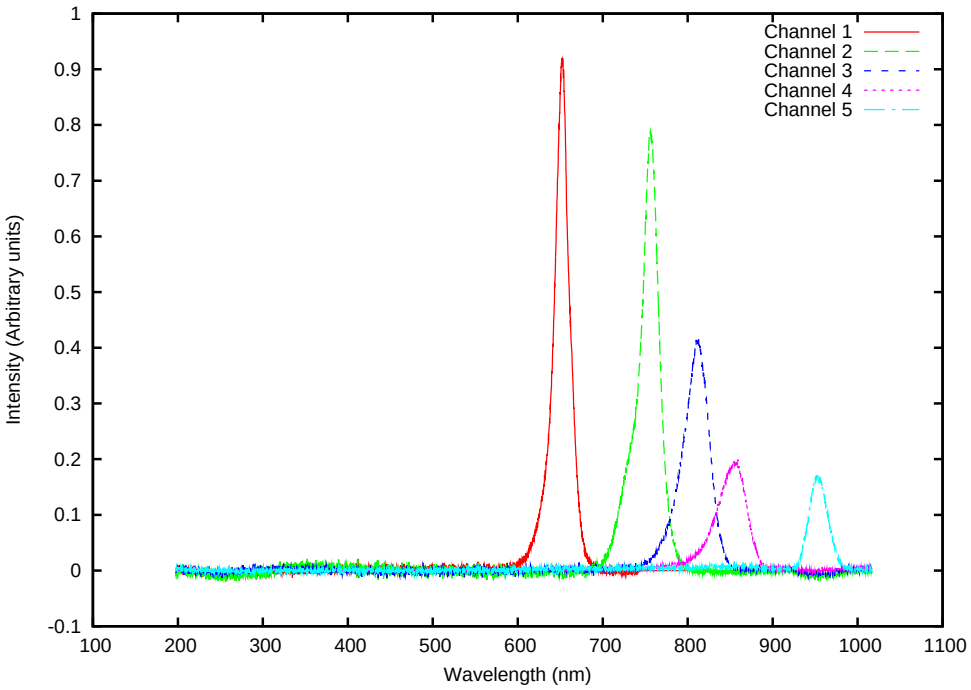


Figure 7.12: The spectrum of the 5 LED sensor

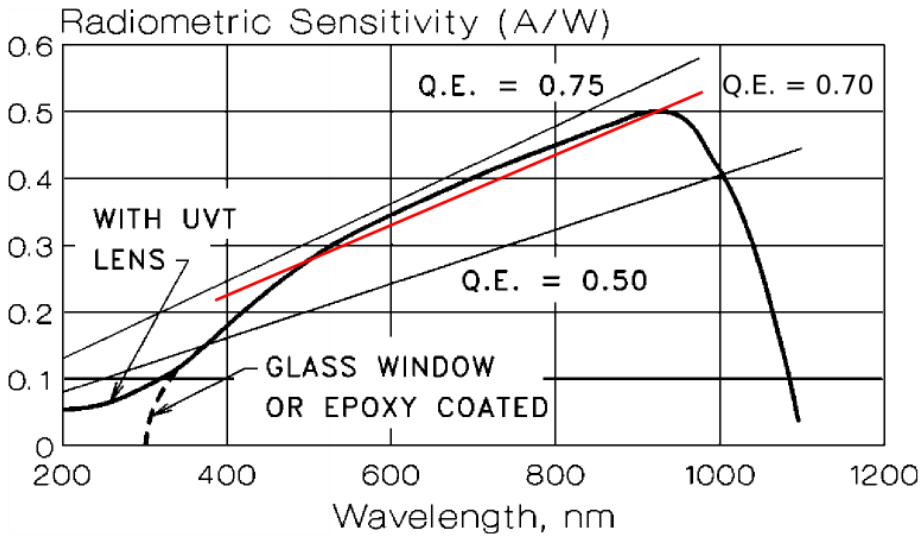


Figure 7.13: The sensitivity of the VTB8441H photodiode used in the sensor PCB. Image © Perkin-Elmer.

sulation (Silbione LSR4340 from Bluestar silicones), reliable adhesion to the PCB and components could not be achieved. Bluestar market a range of adhesion promoting primers and surface treatment agents, but use of these was not attempted due to time constraints and the long lead times of the appropriate primers. As an alternative, the silicone was peeled away from the PCB immediately after it had cured. The PCB and

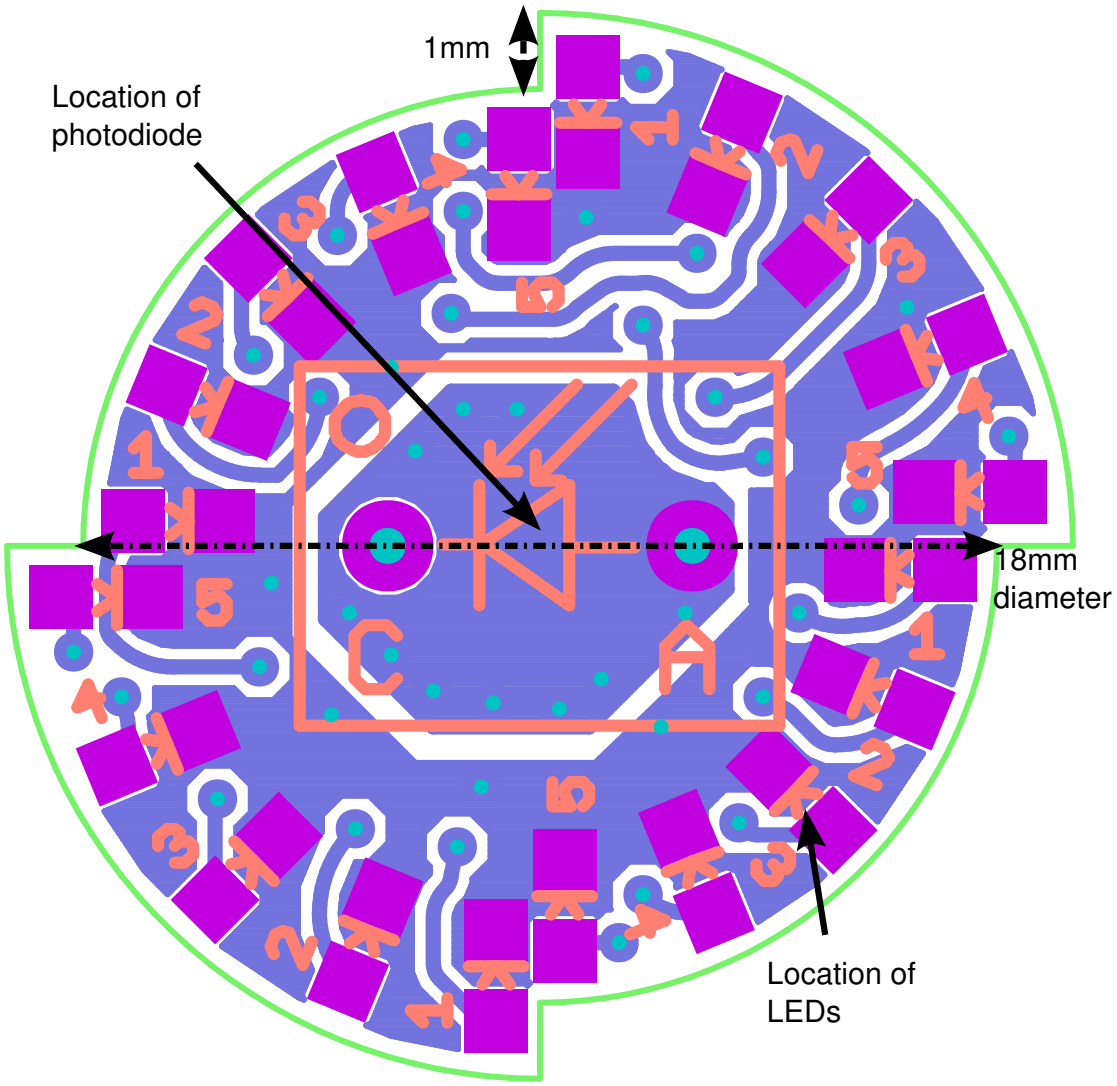


Figure 7.14: The NIRS sensor PCB, a ring of LEDs surrounds the central photodiode.

components were cleaned thoroughly in iso-propanol, before a thin layer of e43 single component room temperature vulcanising silicone sourced from Wacker Inc. was applied to the PCB surface. Finally the molded silicone encapsulation was replaced and the sensor assembly left to cure for several hours. A layer of 6.4mm thick self adhesive closed cell silicone foam sourced from Rogers-Bisco was applied to the back of the sensor to aid even application of pressure.

Figure 7.15 shows the completed sensor with LED cable (white flat flexible cable) and photodiode coaxial cable (grey). The square active area of the photodiode can be seen in the centre. It is clear that there will not be a consistent path length between the different LEDs and the photodiode due to the square active area. There is also likely to be significant light passage horizontally through the transparent silicone, with a significant

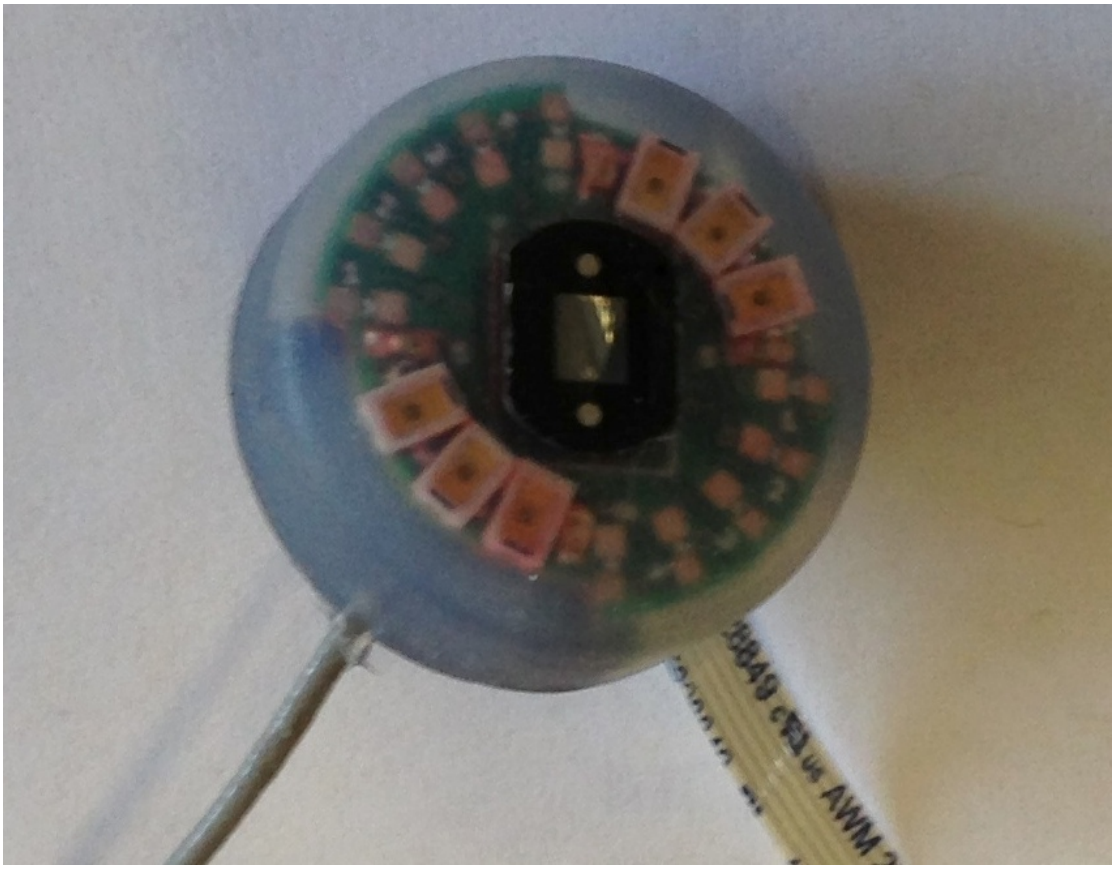


Figure 7.15: The encapsulated NIRS sensor. Ten LEDs are mounted either side of the central photodiode (5 at either side).

fraction of the detected light having entered at the edge of the black ceramic photodiode, especially in the case of heavily attenuating tissue. To eliminate the potential for such problems to occur, a black light absorbent tape with <5% reflection across the near infrared (Thor Labs T743-2.0) was applied to the front face of the detector. A punching jig machined from aluminium was used to form 1mm diameter holes aligned with the LEDs and a 2mm diameter hole over the photodiode. A silicone pressure sensitive transfer adhesive (91022 transfer film from 3M) was used to adhere light shield tape to the sensor as the acrylic based tape backing adhesive on T743 tape is incompatible with silicone rubber surfaces, and has no tack strength. For clarity, figure 7.15 shows the sensor assembly prior to attachment of the light shield.

Refractive index mismatching is a practical problem faced by both optical and ultrasonic medical devices. The refractive index in air is approximately 1, but this increases to approximately 1.4 in biological tissues (for example see Bolin [217]), and a similar value in silicone rubber. Therefore, if an air gap is present between the sensor and skin surface, a significant proportion of the light will be reflected at each interface. From the Fresnel equation under the condition of normal incidence,

$$R = \left| \frac{n_1 - n_2}{n_1 + n_2} \right|^2, \tag{7.28}$$

setting $n_1 = 1.4$ and $n_2 = 1$, $R = 0.27\dot{7}$. It can thus be seen that the intensity loss is not significant enough to cause a serious decrease in signal to noise. However, a signal fluctuation of the order of 20% or more is enough to cause a serious signal artefact if the sensor to skin contact is variable, for example under the effect of variable application pressure.

Glycerol is non toxic, bio-compatible, and widely available, and has a refractive index of 1.47. Its use as a refractive index matching agent for optical studies of tissue has been described by number of authors, for example Tuchin et al [218]. For these reasons Glycerol was used as an index matching agent for the sensor.

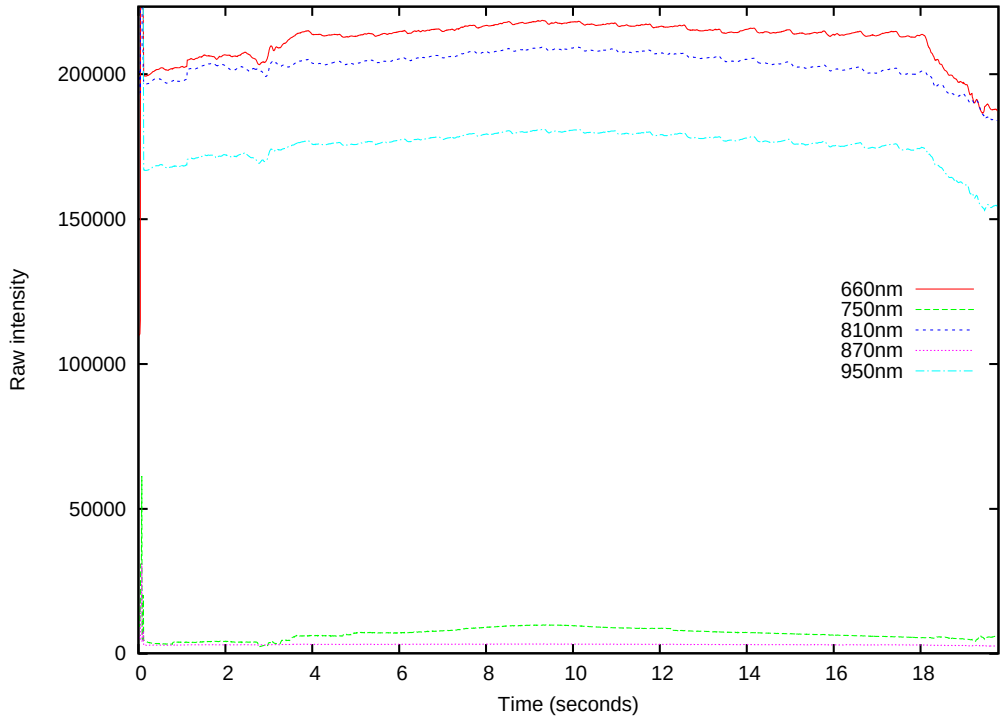


Figure 7.16: Test of the sensor placed on the palm of a healthy volunteer. Two of the channels disabled, showing low crosstalk and a pulsatile PPG.

Figure 7.16 shows data recorded from the palm of a healthy adult volunteer using the five wavelength NIRS sensor. Two of the channels were disabled to allow the inter channel crosstalk to be investigated. The pulsatile signal is clearly visible in the 660nm, 810nm, and 950nm channels as an approximately sawtooth fluctuation 1% modulation depth and frequency of slightly over 1Hz. There is a relatively small fluctuation in the light levels over the duration of the recording which is highly correlated between the multiple channels, suggesting varying application pressure as a likely cause.

The crosstalk into the unused channels is extremely low for the first four seconds of the recording, but then rises considerably, and correlates with the drift in the three modulated channels. This suggests that capacitive coupling contributes significantly to the crosstalk, and so more effective electrical shielding of the LEDs and the flat flexible cable that is used to drive them could significantly reduce crosstalk.

The multi wavelength LED based NIRS sensor is suitable for the study of both CW-NIRS (and thus venous occlusion based measurements of perfusion), and capillary refill, with the 116Hz intensity data output rate being sufficiently high for the transient absorption effects of capillary refill to be observed. Due to the fact that capillary refill based perfusion sensing is a relatively unexplored area in biomedical sensing, it was determined that development of an automated pressure application system might allow useful developments to be made with respect to capillary refill. Furthermore, due to the low resource use of the NIRS system, the "F4 Discovery" development board should be capable of operating a pressure application and control system running as an RTOS task in parallel with the NIRS DSP thread. Development of such a controlled pressure application system is the topic of the next chapter.

Chapter 8

Automation of the Capillary refill test with NIRS and pressure application

As previously discussed in the prior chapter on optical techniques, capillary refill in its existing form is a manual test, use of which is currently widespread in clinical practice. Automation of both the pressure application and the blood volume assessment (through NIRS) offers the potential to improve both the accuracy of the test, and perhaps more importantly allow the origin of the signal and the precise dynamics of the blood flow to be explored. There has been relatively little published quantitative research into the capillary refill test, and no published in-vivo studies involving use of both NIRS and a quickly reduced application pressure could be found. For these reasons the sensor and NIRS system described in the previous chapter was adapted for use in capillary refill studies. This required accurate pressure application to the skin surface with a fast and accurate transient response.

8.1 Development of an automated pressure application system

Use of pneumatic systems is widespread in medical engineering, and a pneumatic system for pressure application would appear to be a suitable route for a commercialised and mass produced capillary refill sensor. However, it would be somewhat challenging to adapt the cylindrical reflection mode NIRS sensor for use with pneumatic pressure application.

A pneumatic cylinder system would not apply a force proportional to the input air pressure due to static friction from the seals. A flexible bladder placed above the sensor would be another option, but in this case the transfer of the tensile stress within the bladder skin to the sensor would be a complex function of both bladder inflation and sensor orientation with respect to the bladder mounting.

8.1.1 Solenoid based capillary refill system

Use of electromechanical actuators was pursued due to the anticipated difficulties associated with using pneumatics. Solenoids have the advantage of high speed, and with only one moving part there is also no gearing to give rise to backlash, making a solenoid based system an obvious choice for use in capillary refill experiments.

An adjustable arm was constructed with a solid base and mounting bolts for the F4-Discovery based NIRS system described in Chapter 6. A spare timer and PWM output channel on the STM32F4 processor being used to drive a solenoid through a MOSFET from a 12V DC supply.

The NIRS sensor was mounted through a Honeywell FSS1500NSB force sensor to allow closed loop pressure control, and an instrumentation amplifier was added to the F4 Discovery daughterboard, passing the amplified transducer voltage to the second ADC (ADC2) on the STM32F4. This allowed a pressure control thread running on the embedded RTOS to measure applied pressure without interrupting the continuous photodiode voltage ADC conversions taking place using ADC1 (the first ADC unit on the micro-controller). Figure 8.1 shows the completed prototype application system. The force transducer was calibrated before use using a test mass designed for use with scales. As the Honeywell datasheet gives linearity as $\pm 0.5\%$, a simple linear voltage response with force was assumed. The validity of this assumption has not been tested, and a more complex multipoint calibration was not attempted, but according to the device documentation all sensors were tested for within bounds performance at the factory. Pressure was calculated from the NIRS sensor diameter of 20mm, i.e.

$$P = F/A = \frac{F}{\pi r^2} = \frac{F}{\pi \times 10^{-4}}. \tag{8.1}$$

Solving this equation leads to:

$$P = 3183F\,(Pa\,units\,to\,4s.f.)\; = 0.4617F\,(PSI\,units\,to\,4s.f.). \tag{8.2}$$

The PWM pulse width to the solenoid drive transistor was used as a means to modulate

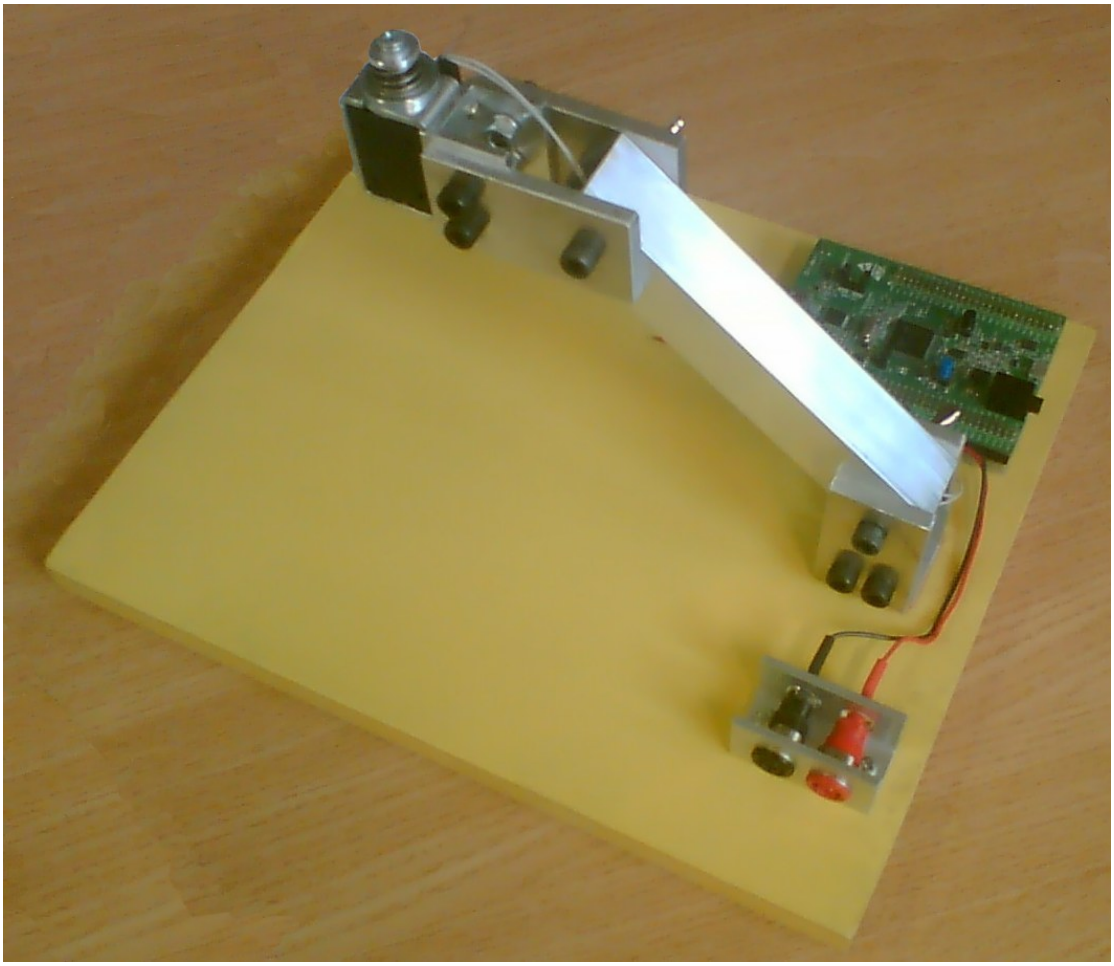


Figure 8.1: The first automated capillary refill device prototype using a solenoid for pressure application. The NIRS system is visible next to the aluminium box section arm.

the applied pressure, and a PID (Proportional, Integral, Derivative) control algorithm was implemented in the pressure control thread. The system timer of the STM32 microcontroller being used to schedule the pressure control loop for a 100Hz iteration rate. Mailbox FIFO buffers were used to pass pressure set-points received over the USB interface by the background thread to the pressure control thread, and also to pass the measured pressure application values from the control loop thread to the NIRS demodulation thread. The 100Hz pressure data were resampled to a 116.2Hz rate inside the NIRS demodulation thread, before being synchronised with the infrared intensity data and passed back to the background thread for broadcast over the USB interface.

The question of how much pressure needs to be applied in the capillary refill test is one that has been discussed extensively in the context of manual assessment. However, definitive guidelines on the necessary application pressure do not exist at the present time. For this reason the maximum likely systolic blood pressure was used as a guideline,

under the assumption that the applied pressure must be sufficient to exclude arterial blood from entering the capillary bed. Taking this as 180mmHg, or 3.5PSI, the maximum force was calculated through application of equation 8.2 as 7.6N.

In practice it was found that the solenoid system had a highly non-linear relationship between applied pressure, armature position, and PWM duty cycle. There was also significant and variable static friction between the armature and the coil. This resulted in all but the least aggressive PID coefficients giving unstable control under some operating conditions. The Ziegler-Nichols tuning algorithm [219] was attempted, but was found to result in PID gains that although stable, gave much slower pressure application and release times than was appropriate for capillary refill studies. Experiments on the palms of healthy volunteers hands with application of Zeigler-Nichols resulted in pressure release times of around 1 second or more for 90% release of pressure.

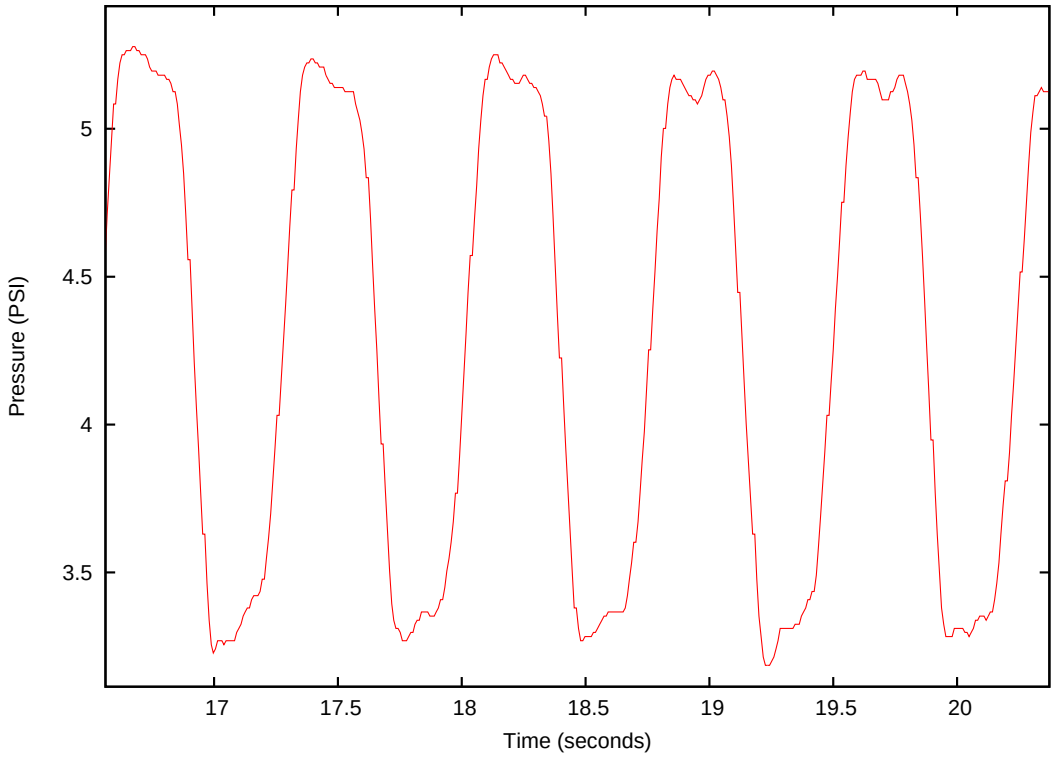


Figure 8.2: Applied sensor pressure in PSI with the solenoid driven using a triangle PWM pulse width function.

Figure 8.2 demonstrates the problems inherent in the solenoid based pressure application system. The PWM duty cycle was varied between 80% and 60% with a triangle function, period 1.5 seconds. It can be seen that the static friction leads to the applied pressure being approximately constant around the peaks of the triangle wave. However, the applied pressure tends to approach the mean pressure during these periods, suggesting that there is also significant visco-elastic behaviour in the volunteer's hand. The

static friction is apparently relatively inconsistent, with the maximum applied pressure changing from peak to peak by some 0.2PSI or thereabouts in some cases. This implies that any control algorithm designed to track the stiction effect would be unable to achieve pressure accuracy of better than $\sim \pm 0.1$PSI, an application pressure that was found to result in significant changes in near infrared attenuation.

The total moving mass (armature + sensor + cabling) was in the region of 120grams. If the sensor is required to move 1cm in 80milliseconds (a time and distance that was judged to be likely to be fast enough to observe any dynamic blood volume effects), then solving for the acceleration and thus the required force gives a peak value of 0.75N. Dividing by the sensor surface area gives an equivalent pressure of 0.34PSI, showing that the inertia of the moving mass can result in a pressure value that is significant when compared to the pressures typically used in capillary refill.

Thus it can be seen that there are numerous obstacles preventing a solenoid system from being accurately controlled. PID control was found to be feasible, but only with an excessively slow response time. An optimal control strategy with modelling of the armature and sensor inertia and the static friction (stiction) may be feasible, but the stiction is relatively inconsistent, and optimal control would not be feasible unless the relationship between solenoid force and PWM duty cycle was known. As the force is also a function of armature position, this is impossible to achieve.

One possible approach, that was not pursued, would be to use a current transformer to monitor the ripple in the solenoid coil current throughout the PWM cycle. The solenoid inductance was found to vary from 58 to 47mH as the armature was moved from fully inserted to fully extended respectably, demonstrating that armature position can be determined from measurement of coil inductance. Determination of the current ramp rate during both parts of the PWM cycle (MOSFET on and off) leads to a system of two simultaneous equations that can be solved to determine both current through the solenoid and its inductance:

$$\frac{dI_{coil}}{dT} = \frac{[V_{supply} - (R_{coil} + R_{MOSFET})I_{coil}]}{L_{coil}} \tag{8.3}$$

as the MOSFET is on, and

$$\frac{dI_{coil}}{dT} = \frac{[-V_{Schottky} - R_{coil}I_{coil}]}{L_{coil}}, \tag{8.4}$$

when the MOSFET is turned off and the current is circulating back to the supply via a Schottky diode. These current ramp rates will be proportional to the current transformer output voltage, and can be measured directly with a spare ADC channel.

Of the terms in these two equations, the supply voltage V_{supply} is known, and R_{coil} can be easily determined, as can R_{MOSFET}, although for most modern power MOSFET transistors this should be negligible. The diode voltage $V_{Schottky}$ can be approximated from the device datasheet and knowledge of the typical current. This leaves only I_{coil} and L_{coil} as unknowns, leaving a system of two equations that can be solved. Bench calibration data for the solenoid could then be used to convert the inductance to an armature position, and calibration curves of force versus I_{coil} and position used to calculate the required iterative change in PWM duty cycle in order to generate an optimal change in armature force. Due to the non negligible armature inertia, a trajectory optimisation algorithm would likely be required for such an optimal control scheme to be achievable. The resistance of the solenoid used on the prototype device was found to be 17Ω, leading to a current decay constant of 3.5 milliseconds. This is a relatively short time relative to the typical duration of solenoid movements, but if very fast movements are desired it may also be necessary to model the effect of coil inductance on current, or alternatively implement closed loop current control using the solenoid current derived from the two current equations.

An alternative open loop current control approach could be used if the highest rate of change of solenoid current was not essential. The equilibrium mean current can be found from equations 8.3 and 8.4 by setting the product of 8.3 with the PWM on (high) time equal to 8.4 with the off (low) time. It is then possible to solve for the equilibrium I_{coil}.

$$I_{coil}(equilibrium) = \frac{R_{coil} + f(R_{coil} + R_{MOSFET})}{fV_{supply} - V_{Schottky}}, \tag{8.5}$$

where $f = T_{on}/T_{off}$, the PWM on time divided by off time.

The problem of stiction would also complicate application of such a technique in practice. However, induced vibration is a widely used technique to mitigate the effects of stiction, and small oscillations of the solenoid at frequencies above the region of interest for NIRS studies (i.e. in the region of >30Hz) may allow stiction to be minimized without significant corruption of the NIRS data.

Although it can be seen that there is the potential for a solenoid based pressure control system to achieve acceptable accuracy and responsiveness, it was determined that a linear actuator system had more potential in a short development time. For this reason the solenoid system was abandoned, but it may have the potential to provide a more cost effective or even faster and more accurate pressure control if the control technique that has been outlined was perfected.

8.1.2 Linear actuator based capillary refill system

Unlike a solenoid system where only the instantaneous force on the armature can be controlled, a linear actuator allows the sensor position to be changed directly via the lead screw drive motor. Actuator position can then be set using a relatively simple control loop such as PID. Alternatively, use of an off-the-shelf linear actuator system would mean that only a control loop to regulate set-point position based on the measured application pressure would be required.

There are a large number of commercial off-the-shelf linear actuators, but most are designed for industrial applications and are either too large, costly or require complex power supply arrangements. For this reason, rather than a high performance industrial system, use of a Firgelli L12 actuator (Firgelli Technologies Inc. Victoria, BC, Canada) was attempted as a solenoid replacement. This is a small linear actuator designed for operation with 5V remote control systems, and was specified for 12N static force and 25mm/s speed. The 25mm/s maximum speed was found to be too low for fast pressure control to be feasible, and a typical pressure release from 3.5PSI to 0 took around 200ms.

Increasing the operating voltage in an effort to reduce pressure release time resulted in unstable position control by the actuator's built in motor controller. Thus in order to achieve a fast response time with a combination of accurate and stable positioning, both the geared DC motor and the control PCB were removed from the Firgelli actuator and replaced by a NEMA11 stepper motor (Astrosyn Ltd, Kent, England) and external stepper motor controller (a breakout board for the Allegro A4983 microstepping motor driver). The output of the membrane potentiometer built into the Firgelli actuator was used as a means of coarse position measurement, with integration of the number of steps being used for precise positioning.

Figure 8.3 shows a diagrammatical representation of the experimental setup. The stepper motor was controlled from the existing pressure control loop thread running at 100Hz on the micro-controller. The MY3002 stepper motor has a relatively high holding torque of up to 90mNm, which was found to be sufficient for an actuator thrust of over 40N with the Firgelli actuator lead-screw. During periods of high acceleration the stepper motor moment of inertia may become significant. The MY3002 series motors have a moment of inertia of $1.2 \times 10^{-6} Kgm^2$, meaning that, as the Firgelli lead-screw pitch is 4mm, a linear acceleration of $6ms^{-2}$ corresponds to an angular acceleration of over 9400 radians per second squared. This gives rise to an effective torque loss of 11mNm. Thus it can be seen that the maximum likely actuator acceleration can be achieved without a large torque loss.

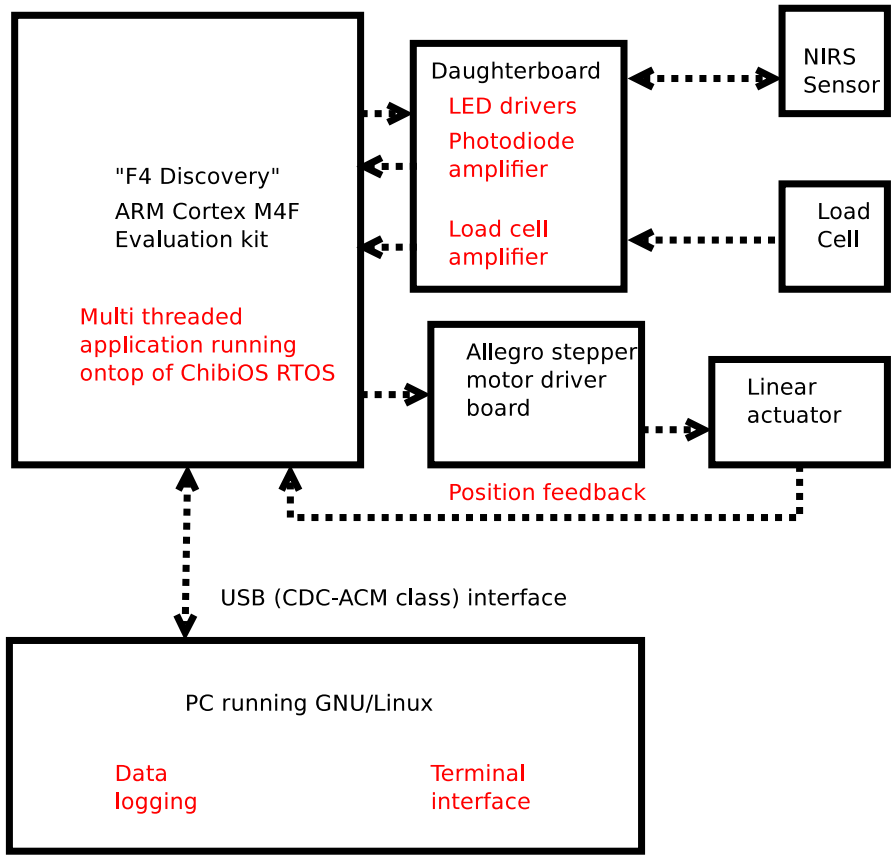


Figure 8.3: System diagram of the linear actuator based pressure controller.

However, the stepping rate of the motor must not be changed too abruptly, or there will be an extremely high angular acceleration transient that may cause the motor to lose steps. The "pull in torque" is a parameter listed for most stepper motors. This is defined as the maximum torque that can be applied to the motor shaft during an instantaneous change in stepping rate without missing steps. The MY3002 pull in torque is approximately constant at 3Ncm between zero and 200 steps per second. Above 200 steps per second the pull in torque drops rapidly to zero at 500 steps per second - meaning missed steps if the stepping rate is changed abruptly by more than 500 steps per second, even with no applied load. For this reason a trajectory optimisation algorithm was used to generate a series of motor step rates at 400Hz iteration rate that would reach the target actuator position without exceeding a pre-set acceleration limit and with no overshoot. A timer interrupt synchronised to the 100Hz control loop thread and triggering at 400Hz was used to set the PWM periods for the stepper motor drive signals. Considering the 9400 radians per second squared angular acceleration discussed earlier, it can be seen that this corresponds to a 3×10^5 steps per second squared step rate acceleration, or 750 steps per second maximum change in step rate between two consecutive 400Hz periods. This exceeds the pull in torque of the motor, so for this

reason the pre-set acceleration limit was set to a lower value of $3ms^{-2}$. This resulted in a slight decrease in the rate of change of applied pressure. Experiments on the hands and forearms of healthy volunteers indicate that a 0 to 3.5PSI applied pressure change requires an actuator movement of approximately 1cm maximum. At $3ms^{-2}$ acceleration this takes 115ms, which is moderately significant compared to the shortest capillary refill times typically encountered, which are of order 800ms.

In practice it was found that the stress versus strain relationship for sensor pressure versus position was linear to a reasonable approximation during tests on volunteers. Visco-elastic effects were also present, but the relatively slow decay times for visco-elasticity (around 1 second or so) compared to the actuator response time meant that ignoring these effects and simply adjusting actuator position to account for both movement of the volunteer and visco-elasticity was found to yield acceptable performance.

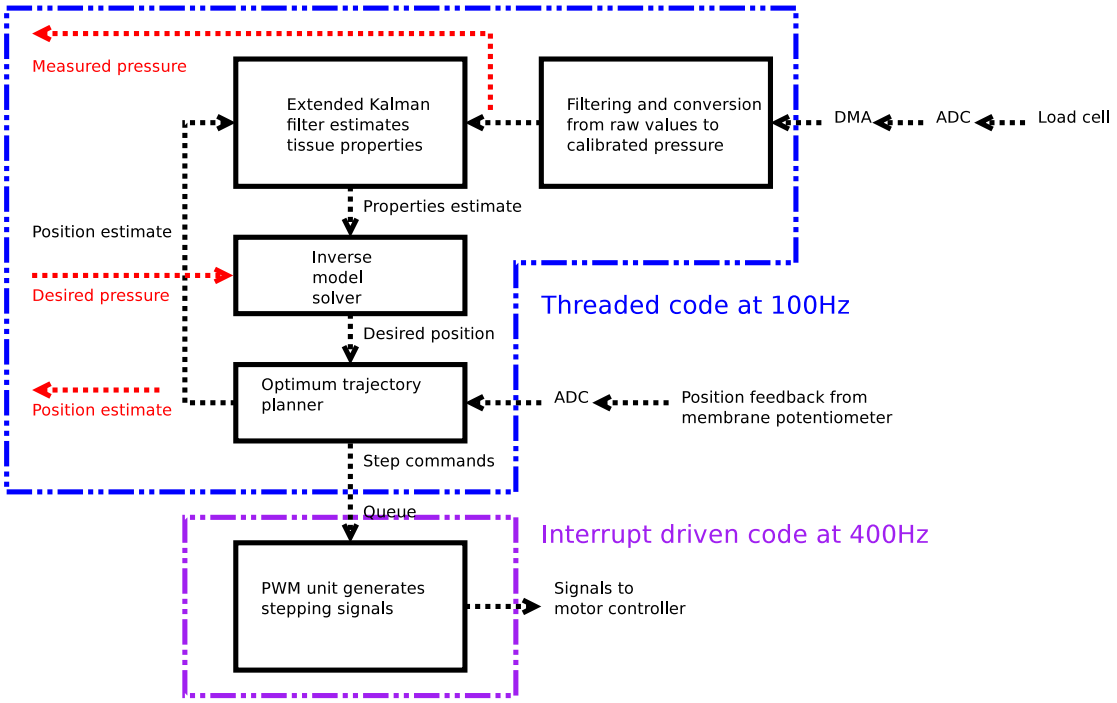


Figure 8.4: Diagram of the pressure control system

Figure 8.4 shows the final overall form of the pressure control system. It was found that the stress/strain relationship became less linear at higher applied pressures (>2PSI), and was variable between volunteers. This is similar to the finding of Hendriks et al [220].

In order to allow the device to be used with multiple volunteers without calibration, an Extended Kalman filter was used to fit a second order elasticity model to the data in real time

$$\frac{\partial P}{\partial x} = k_0 + k_1 P. \tag{8.6}$$

The two elasticity coefficients (in equation 8.6) were combined with a pressure estimate to form the Kalman filter's state vector:

$$X = \begin{pmatrix} k_0 \\ k_1 \\ P \end{pmatrix}. \tag{8.7}$$

Linearisation of equation 8.6 for small position changes leads to:

$$\Delta P = (k_0 + k_1 P)\Delta x. \tag{8.8}$$

Which was used together with the known actuator positions as the prediction step in the filter. The update stage of the filter used the measured pressure to update the state P. Thus the entire filter can be described as:

$$X_{n|n-1} = f(X_{n-1|n-1}, u_{n-1}), \tag{8.9}$$

where f is the update model (8.8) and u the control input, i.e. Δx

$$P_{n|n-1} = F_{n-1} P_{n-1|n-1} F_{n-1}^T + Q_{n-1}, \tag{8.10}$$

where Q is the covariance matrix for the accumulating random errors in the state vector.

$$y_n = P_{measured_n} - \begin{pmatrix} 0 & 0 & 1 \end{pmatrix} . X_{n|n-1} \tag{8.11}$$

$$S_n = H_n P_{n|n-1} H_n^T + R_n, \tag{8.12}$$

where R will be the error variance in the pressure measurement supplied to the filter.

$$K_n = P_{n|n-1} H_n^T S_n^{-1}, \tag{8.13}$$

$$X_{n|n} = X_{n|n-1} + K_n y_n, \tag{8.14}$$

$$P_{n|n} = (I - K_n H_n) P_{n|n-1}, \tag{8.15}$$

where H is the Jacobian matrix for the measurement model, which will simply be $\begin{pmatrix} 0 & 0 & 1 \end{pmatrix}$.

The Jacobian matrix for the update model, F, has a more complex form derived from the partial differentiation of equation 8.8:

$$F_{n-1} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \Delta x & P_{measured}\Delta x & 1 + X_{n|n-1}(2)\Delta x \end{pmatrix}. \tag{8.16}$$

The coefficient R was pre-set to an appropriate value using pressure sensor data analysed in GNU-Octave to determine the variance. The matrix Q was diagonal, i.e. it was assumed there would be no covariance between errors in any of the state vector components. It was assumed that the coefficients k_0 and k_1 would remain constant throughout the observation period, so the middle and upper left components of the Q matrix diagonal were set to very small values. However, the initialisation covariance for P_0 had large values for these components, as the initial physical properties of the tissue under test were completely unknown. The third component of the Q diagonal vector represents the increasing variance for the state vector P, and was assumed to represent the accumulation of mis-modelling errors and the drift in hand position. An arbitrary large value was used here to allow hand movements to be accommodated in a stable manner.

In order to investigate the stability of the filter under varying initial conditions, inputs and pre-set noise covariance matrices, a simulation was written of the Kalman filter, actuator system, and an approximated area of tissue (non-linear behaviour was assumed with no visco-elastic or thixotropic components). As this was written in the C programming language, and used the same matrix libraries cross compiled to the x86 architecture rather than ARM. This allowed the entire control system code to be debugged in a simulated environment on a PC, so allowing carefully controlled and repeatable tests without risk of damage to hardware. The optimal values for the Q matrix diagonal and the initialisation of the state covariance matrix P_0 were determined via this method.

The ARM CMSIS™matrix library was used for the matrix operations. This library uses the FPU acceleration on the Cortex M4F based STM32F4, giving single cycle single precision (32 bit) multiplication and addition for the large number of matrix operations involved in each iteration of the filter.

The Extended Kalman filter provides a fitted model of tissue properties; a set of coefficient which can be combined with the current pressure set-point to yield an estimated

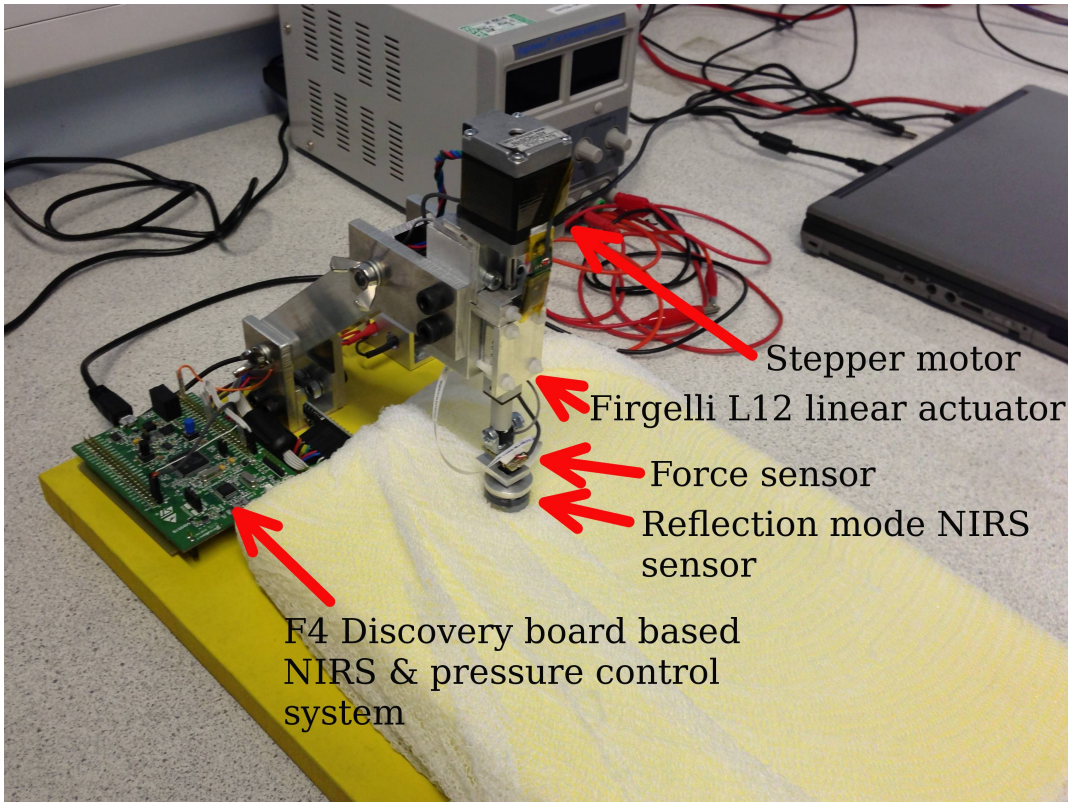


Figure 8.5: The automated capillary refill test rig with linear actuator fitted.

actuator position at which the applied pressure and set-point pressure will be equal. Returning to equation 8.6, and integrating with respect to x:

$$\Delta P_n = \int_{x_0}^{x_n} k_0 + k_1 P \, dx. \tag{8.17}$$

This has a solution:

$$P_n = Cexp(k_{1_n} x) - k_{0_n}/k_{1_n}, \tag{8.18}$$

where C is a constant setting the initial pressure at the reference point $x = 0$. If the reference point is continually redefined as the current position with each iteration of the filter, then introducing set-point pressure S_n this becomes:

$$S_n - P_{measured_n} = C_n \left[exp(k_{1_n} \Delta x_{target}) - 1 \right], \tag{8.19}$$

$$P_{measured_n} + k_{0_n}/k_{1_n} = Cexp(k_{1_n} x). \tag{8.20}$$

As the distance x is continually redefined as zero, this becomes:

$$P_{measured_n} + k_{0_n}/k_{1_n} = C. \tag{8.21}$$

Equation 8.19 can now be solved to determine the optimal movement distance Δx_{target}:

$$\frac{ln\left[\frac{S_n - P_{measured_n}}{P_{measured_n} + k_{0_n}/k_{1_n}} + 1\right]}{k_{1_n}} = \Delta x_{target}. \tag{8.22}$$

This function was implemented using an FPU accelerated logf function, which resulted in low clock cycle use per iteration on the STM32F4 (actual number of cycles fluctuated but was <500 for all the data tested).

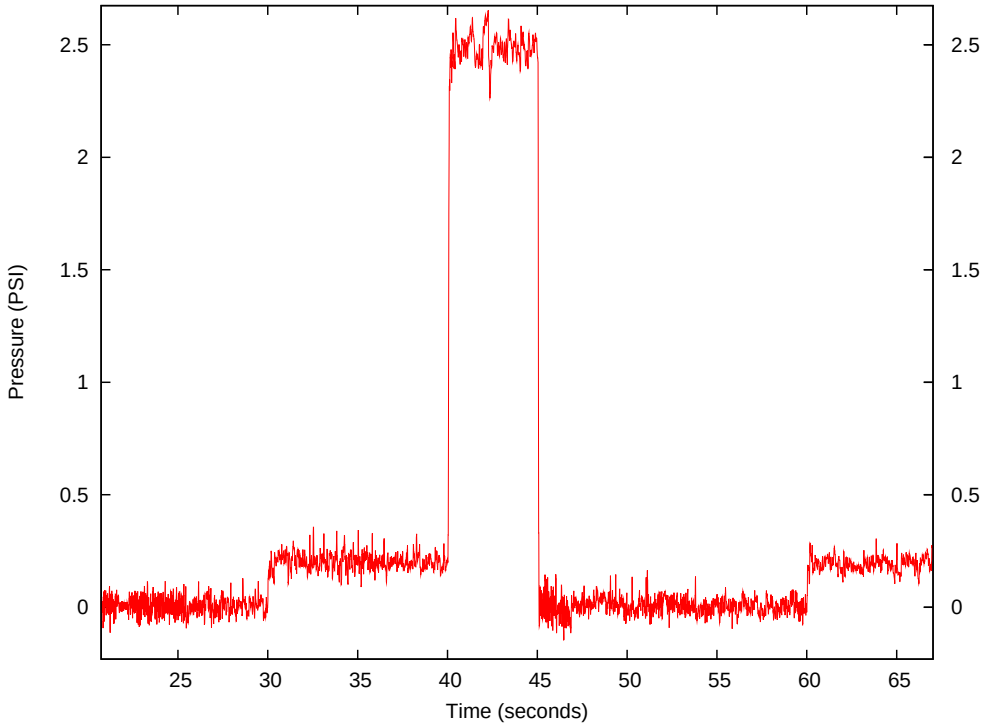


Figure 8.6: Pressure versus time from the linear actuator system applied to the forearm of a healthy volunteer

Figure 8.6 shows the control system implementing the optimal control algorithm to apply pressure to the forearm of a healthy volunteer. The pressure set-point was changed from 0 to 0.2 to 2.5PSI, and the applied pressure can be seen to track the set-point with a very fast response. Rise time from 0.2 to 2.5PSI is in the order of 80 milliseconds.

It can be seen that there is significant noise on the applied pressure values. The Honeywell force transducer produces a very weak (millivolt level) differential output voltage, which is vulnerable to both voltage ripple at the instrumentation amplifier and electromagnetic interference. The noise was found to be mainly due to electrical interference from the stepper motor. In order to mitigate the effect of this interference, the ADC2 sample rate was increased to 14kHz, and the quick-sort algorithm (an efficient sorting algorithm that is on average O(nlog(n)) [221]) was used in order to discard the upper and lower quartile from the data. The remaining data were then averaged to yield a

more stable estimate of the applied pressure. This process is illustrated as part of the data flow in figure 8.4, but was not applied when the data in figure 8.6 were recorded.

Ferrite beads were attached to all stepper motor control cabling between the motor driver board and the F4 Discovery board to further reduce electromagnetic interference.

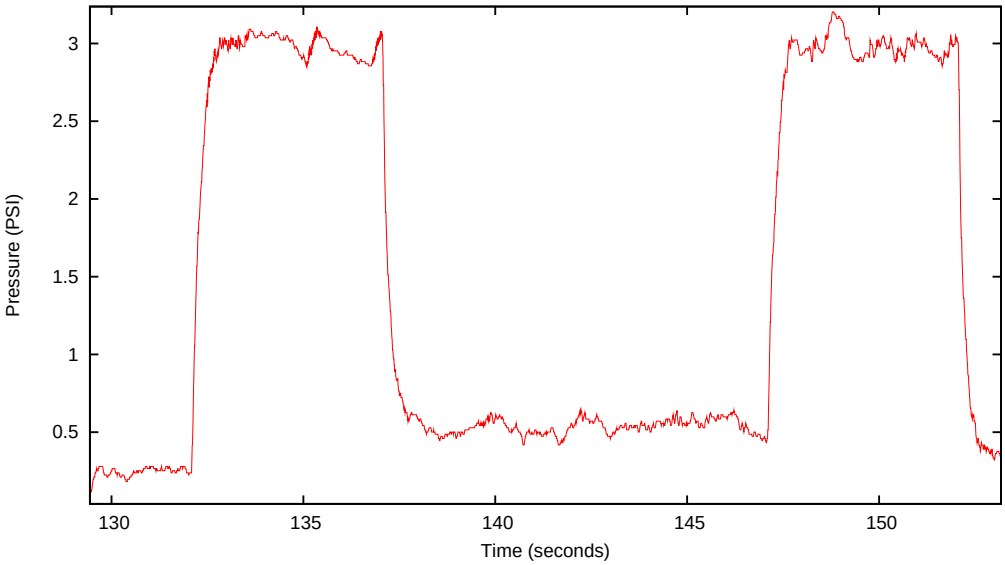


Figure 8.7: Pressure versus time from the linear actuator system applied to the forearm of a healthy volunteer, mean quartile filtering applied

Figure 8.7 shows a plot of a typical pressure versus time sequence with the NIRS sensor applied to the forearm of a healthy volunteer with the optimal control and pressure filtering techniques show in figure 8.4 applied. It can be seen that the filtering results in a clear reduction in the high frequency noise compared to figure 8.6. There is still significant drift over longer periods, with an apparent 1/f characteristic. This was found to be due to small involuntary muscle movements in the arm, and the white polyurethane foam cushion visible in figure 8.5 was introduced to allow the arm to be held with a more relaxed posture in order to minimize this effect.

8.2 Analysis of NIRS data from automated capillary refill

The custom made NIRS system used with the fully automated capillary refill test equipment only provides raw optical intensity levels. To be able to gain an insight into the dynamics of the micro-circulation in a pressure affected tissue volume, these need to be processed into a more usable form. Ideally tissue haemoglobin and oxy-haemoglobin molarities would be obtained for each set of intensity values.

8.2.1 Pulsatile modulation depth fitting

Although a relatively simple differential path length based model would be appropriate for the reflection mode NIRS sensor, a more speculative approach based upon a pulastile amplitude fitting technique was taken initially. With measurements of the pulsatile modulated depth at five separate wavelengths, there are sufficient degrees of freedom to apply the model used in the previous optical methods chapter. By measurement of the modulation depth between pressure applications, when the tissue is in the unperturbed "baseline" state, application of the fitting algorithm yields a solution for baseline tissue haemoglobin and oxy-haemoglobin. With known baseline parameters, the diffusion approximation can be applied to convert raw intensities and generate a set of values for tissue μ_a change during the pressure application and release periods. A technique such as linear least squares (i.e. multiplication by the Moore-Penrose pseudo-inverse of the absorbances matrix) can then be used to convert these absorption changes back into changes in tissue Hb, HbO_2, and H_2O molarities, as well as a fit error metric. By assumption that change in H_2O content indicates transport of interstitial fluid, it may be feasible to scale the absolute value of the tissue scattering coefficient $\mu'_{s(tissue)}$, improving the accuracy of the results over significant pressure changes.

In practice, measurement of modulation depth is not trivial. Simple band pass or comb filtering can be used to remove the majority of noise affecting the signal, but heartrate is variable, leading to a requirement for significant filter bandwidth. To apply filtering the mean heartrate must also be determined accurately. This is itself a non-trivial process, as the capillary refill signal is extremely large relative to the pulsatile component. This means that in order for the heart rate to be observable in the frequency spectrum, the pressure application, release, and capillary refill periods must be excluded from the data. To provide the maximum quantity of capillary refill data, it is desirable to apply pressure as often as practical, leaving the minimum possible time for the pulsatile fluctuations to be recorded between the pressure applications and clear of the period of likely refilling blood. This means that both the signal to noise ratio and frequency resolution that can be obtained from a single interval of pulsatile data are limited.

If multiple observation periods could be combined, this would yield a high signal to noise ratio and high frequency resolution peak at the frequency of the heartrate. The Lomb-Scargle periodogram is one widely used method to obtain power spectral density of gappy and unevenly sampled data. Roberts et al [222] presented an algorithm dubbed "CLEAN" which is also relatively well used. CLEAN has been found to outperform Lomb-Scargle in some circumstance, with fewer sampling time induced artefacts appearing in the output (for example by Baisch and Bokelmann [223]). Use of both of

these methods was attempted, and both were found to be capable of correctly identifying the heartrate peak in the power spectrum.

It was found that application of polynomial de-trending with an order equal to the number of refills in the dataset minus two performed well as a means of minimizing the longer term motion artefacts and drifts in illumination intensity, and gave improved accuracy in the heart rate determination, particularly with challenging low modulation depth recordings. Due to the effective removal of motion and drift, de-trending was applied as a first processing stage to the data used in both modulation depth determination and NIRS fitting in the capillary refill and pressure application regions (this was possible as the de-trend fit polynomial extended throughout the entire dataset).

With an estimate of the heartrate, it is possible to create a comb filter to select only the pulsatile signal and its harmonics. A fixed bandwidth of $f/5$ was used, where f is the principal harmonic frequency (i.e. heartrate(bpm)/60). A second order low pass filter with a transition frequency of 12Hz was placed after the comb filter to attenuate the higher harmonics, where there was likely to be a lower signal to noise ratio (i.e. lower pulsatile harmonic amplitude over noise amplitude ratio). As the absolute amplitude of the pulsatile component is not required, it is possible to use correlation techniques to determine the relative amplitude of the pulsatile signal from the cross correlation between optical channels of the comb filtered data. This can be illustrated in the simple case of three optical channels as follows:

$$\frac{P_2}{P_1} = \frac{Corr(S_2, S_3)}{Corr(S_1, S_3)}, \tag{8.23}$$

$$\frac{P_3}{P_1} = \frac{Corr(S_2, S_3)}{Corr(S_1, S_2)}, \tag{8.24}$$

where $P_n\,(n \in 1, 2, 3)$ is the pulsatile signal, and $S_n\,(n \in 1, 2, 3)$ the filtered optical signal.

Figure 8.8 illustrates the data flow in the automated modulation depth processing software. This was implemented in the GNU-Octave language.

The modulation depths output from this stage of the data processing are measured across a set of relatively broad spectral peaks (see figure 7.12). This makes the point spectrum assumptions employed in the earlier work analysing published modulation depth data with a 10nm line width monochromator invalid. An analysis of the equations developed to describe the pulsatile modulation depths (in particular equation 6.20) shows that the modulation depth has a linear dependence on blood absorbers, and an inverse square root dependence on tissue absorbers. It can therefore be seen that two sets of effective absorbency are needed for each optical channel and absorber combination,

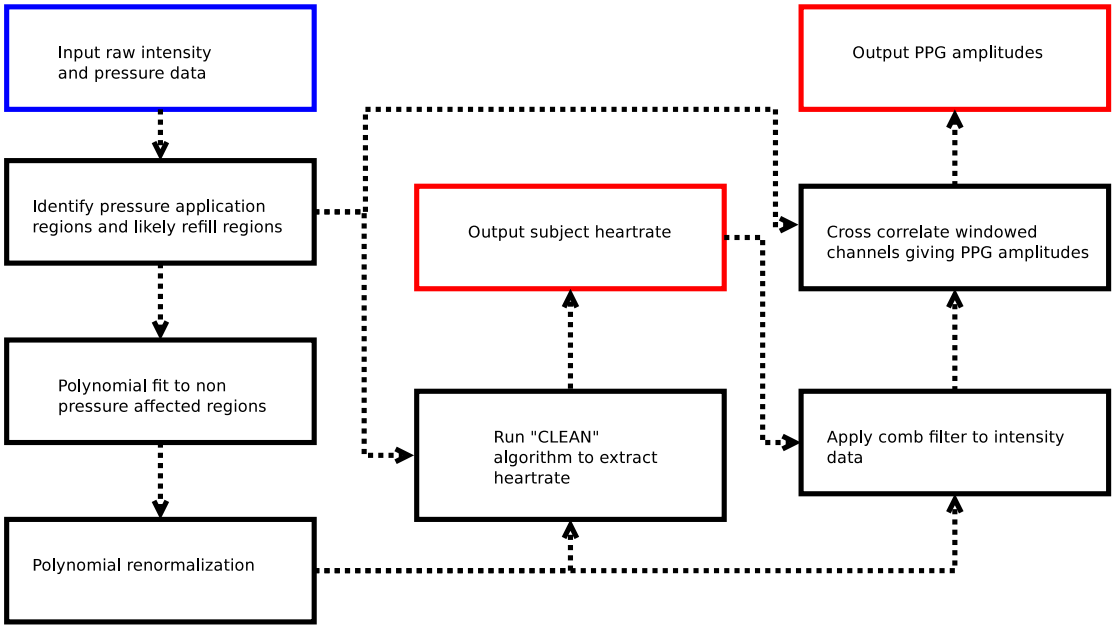


Figure 8.8: Diagram of dataflow in automated capillary refill NIRS processing

one proportional to μ_a integrated over the channels emission function, and the other to $\sqrt{\mu_a}$.

In order to accurately calculate these new sets of effective absorption coefficients, both the calibrated spectrometer data and the sets of raw absorption coefficients (taken from Prahl [224]) were re-sampled to 0.1nm resolution using cubic spline interpolation. Photodiode A/W response was calculated assuming a uniform quantum efficiency of 0.7 (see figure 7.13). Although absolute photodiode sensitivity is not relevant to the absorbency calculation, the change in A/W response with wavelength must be accounted for when calculating the effective emission function for each channel. Finally, the integral of absorbencies multiplied with normalized channel emission function across wavelength was calculated with a cumulative trapezoid method.

Figure 8.9 shows a plot of the NIRS modulation depth recorded from the palm of a healthy volunteer's hand, the positive vertical axis being proportional to attenuation (defined as $A = \ln(I_t/I_{mean})$). On close inspection, some heart rate variability is evident from the slight variance in the peak spacing. The dicrotic notch (a small drop in aortic pressure coinciding with the closure of the semilunar valves and preceding the dicrotic wave) is also visible as a small fluctuation in intensity between the major peaks. There are clear beat to beat fluctuations in the precise shape and timing of this fluctuation, but this is likely to be mostly due to noise as the amplitude of the fluctuations were of the same order of magnitude as the typical sensor noise. With the NIRS sensor placed on a block of absorbent foam, the RMS modulation depth from noise with the comb filter set to 0.75Hz was equal to approximately 10^{-4}. The precise level of noise was

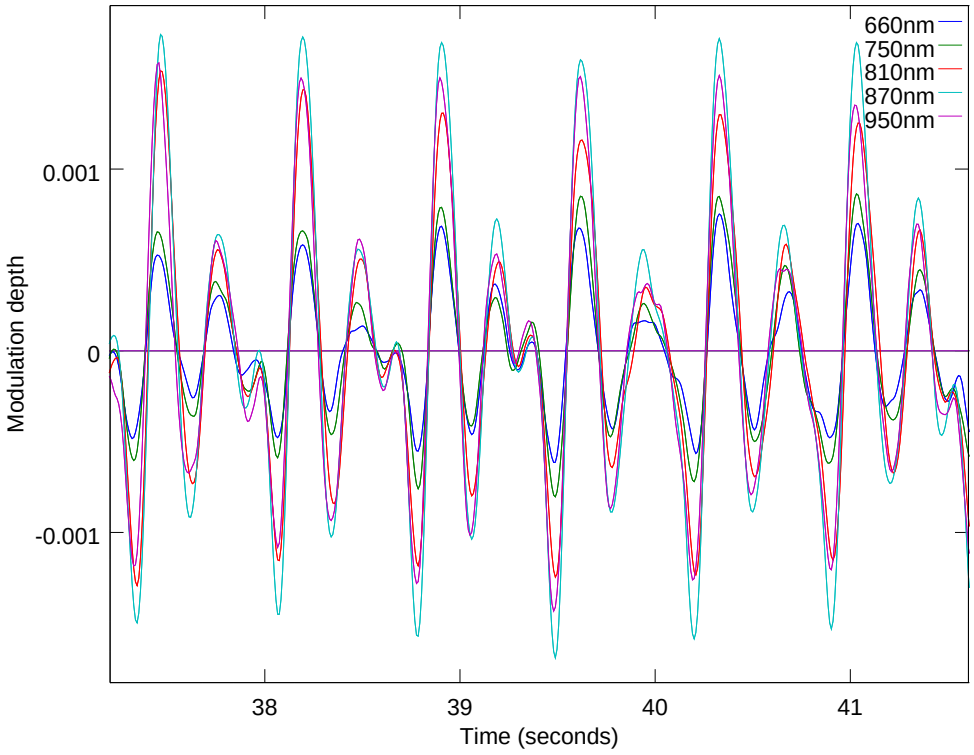


Figure 8.9: The filtered pulsatile signal at the output from the comb filter. Filter output divided by the polynomial fit to give absolute modulation depth values.

found to fluctuate with time, LED intensity and the tasks running on the STM32F4 processor. Measuring the voltage noise on the LED supply rail confirmed it to be a significant source of the observed noise, and as this supply rail was taken from the same 3.3v digital rail used to power the STM32F4 on the discovery board it can be seen that there is significant potential for interference.

However, making design alterations in an effort to improve the stability of the LED drive currents was judged to be unnecessary, as the relative modulation depths can be easily determined even with a low 0.15% modulation depth as in figure 8.9.

Figure 8.10 shows a typical recording of capillary refill sequences on the palm of a healthy volunteer's hand. Target pressure was changed between a baseline level and a 3PSI tissue blanching level over a period of 15 seconds, with a 5 second blanching period (matching the application time typically used in clinical studies) the baseline level was alternated between 0.2 and 0.5 PSI. Actuator position (measured in the extension direction) is plotted in mm units using the same axis as the applied pressure data. It can be seen that there is a hand movement of approximately 4mm at 75 seconds, but the control

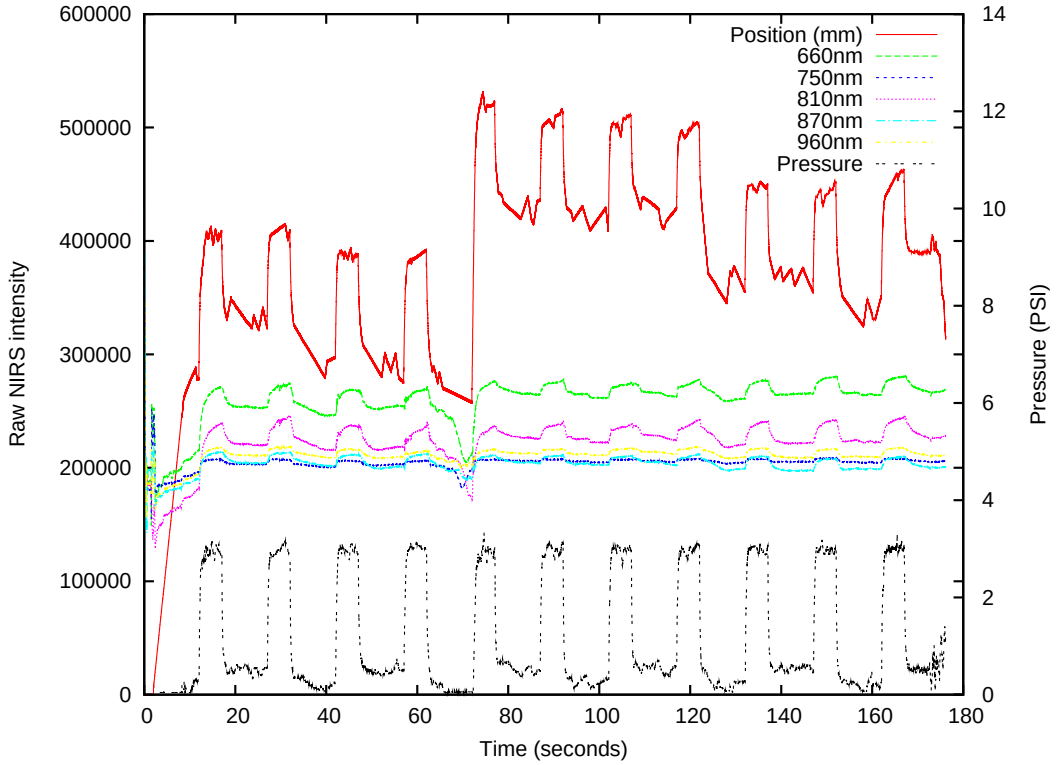


Figure 8.10: Raw NIRS intensity data plotted with applied pressure and actuator position during a series of capillary refill test on the palm of a healthy volunteer's hand.

loop responds quickly, and the subsequent refill does not have degraded accuracy as a result of the unpredictable motion.

For the first ten seconds there is no applied pressure as the sensor has not yet made contact with the skin surface. The actuator extends relatively slowly, much slower than its motion during pressure application and release, as the update stage of the extended Kalman filter was made to assign a much greater noise covariance R if there was no sensor contact with the skin. In the case of the sensor not making contact with the skin surface, there is still force sensor noise giving fluctuations around zero, so a threshold of 0.05PSI was set to determine if contact had occurred.

As well as giving a smooth initial application, this change in the R value can be seen to have mathematical validity, as the non-linear stress versus strain model breaks down in the case of a sensor that is not in skin contact, so assigning a large error to the measurements is one highly approximated way to handle this.

Significant extension and withdrawal of the actuator can be seen in the periods between the pressure applications. This was found to be due to the significant backlash in the Firgelli actuator mechanism. This can be seen to be equivalent to a positioning error

of approximately ± 0.15mm. Careful observation of the pressure perturbations reveals that a significant number of the fluctuations occur simultaneously with the stepper motor reversing direction as the lead-screw engaged inside the actuator. Thus it can be seen that actuator backlash is a significant source of applied pressure noise, and steps to minimize backlash (such as using a higher tolerance actuator) would lead to a significant improvement in the data compared to the Firgelli L12 actuator.

Following use of the modulation depth algorithm, the data inversion used the following system of equations:

$$I = I_0 exp(-\mu_{eff} d). \tag{8.25}$$

Following the diffusion model approximation developed by Farrell and Patterson [215], tissue attenuation was modelled using an effective attenuation

$$\mu_{eff} = \sqrt{3\mu_a(\mu_a + \mu_s')}. \tag{8.26}$$

As $\mu_s' \gg \mu_a$, this can be approximated as:

$$\mu_{eff} = \sqrt{3\mu_a \mu_s'}. \tag{8.27}$$

Introducing a reference intensity I_{ref}:

$$I_{ref} = I_0 exp(d\sqrt{3\mu_{a_{ref}}\mu_s'}), \tag{8.28}$$

introducing the equivalent equation for the measured intensity values I_n, this can be rearranged to form:

$$\frac{ln\left[\frac{I_n}{I_{ref}}\right]}{d} = \sqrt{3\mu_{a_{ref}}\mu_s'} - \sqrt{3\mu_{a_n}\mu_s'}, \tag{8.29}$$

$$\frac{\left[\sqrt{3\mu_{a_{ref}}\mu_s'} - \frac{ln\left[\frac{I_n}{I_{ref}}\right]}{d}\right]^2}{3\mu_s'} = \mu_{a_n}. \tag{8.30}$$

It can be seen from the denominator that in order for changes in attenuation from the baseline level to be converted to absolute molarity changes, the tissue scattering coefficient μ_s' must be known in absolute units. Additionally, there is a non-linear dependence due to the μ_s' term inside the numerator expression, meaning that accurate measurement of μ_a changes in arbitrary units also require μ_s' to be known. For this reason an estimate was made for the coefficient k in equation 6.17, giving an estimated $\mu_s'(\lambda)$ for use in equation 8.30. The k parameter was set such that the transport corrected scattering coefficient at 800nm was equal to $0.45mm^{-1}$, based on the results presented by Farrell et al [215].

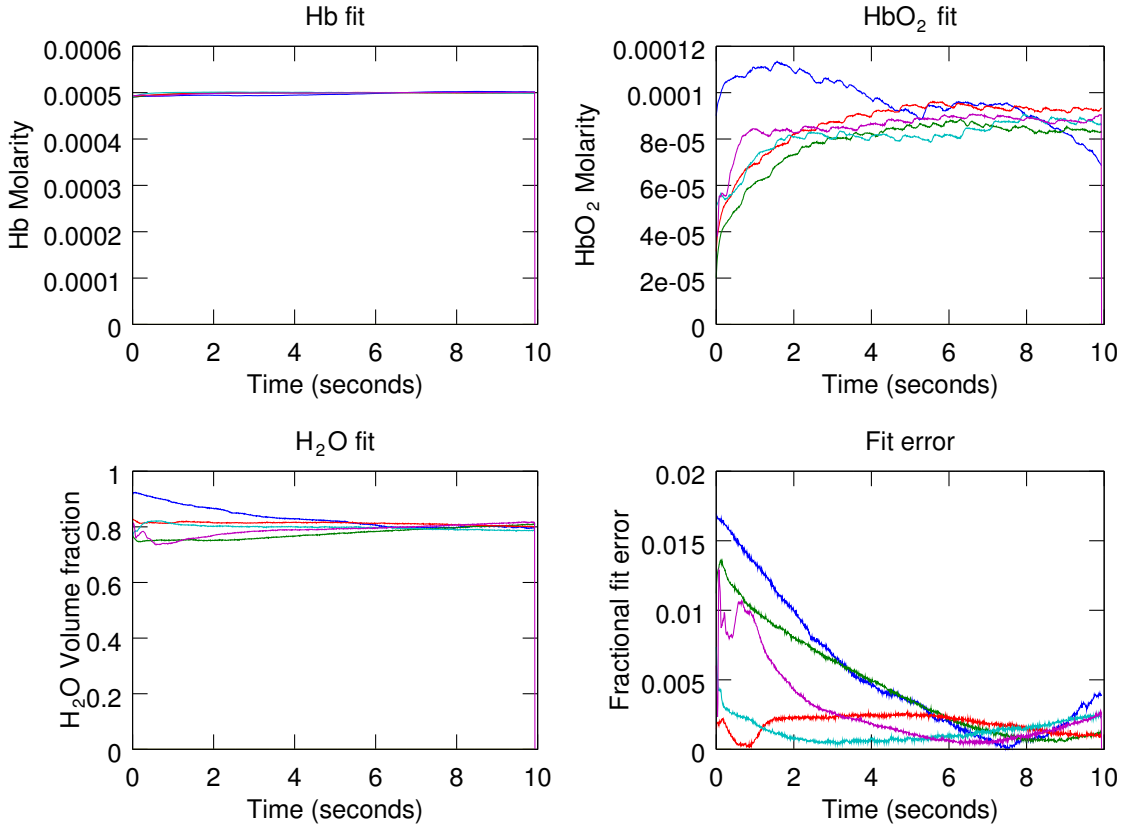


Figure 8.11: NIRS capillary refill data for multiple refills processed using pulsatile modulation depth fitting. Data from healthy adult volunteer, each line is a single refill.

Figure 8.11 demonstrates inversion of capillary refill data recorded from the palm of a volunteer's hand to obtain total molarities of haemoglobin, oxy-haemoglobin, and water versus time after pressure release from 3PSI to 0.2PSI. The pressure was not released to zero, as this was found to give a risk of the sensor partially lifting from the skin surface, resulting in large coupling artefacts due to reflection at the skin surface and changing effective path length through the tissue.

The inverted data span the entire ten second time interval between pressure release and the start of the pressure application for the next refill, and five second interval preceding each pressure application was used to obtain the pulsatile modulation depth. The fit error is also plotted, this was defined as the length of the $\mu_{a_n}(err)$ error vector divided by the length of the μ_{a_n} vector, where the error vector was defined as the difference between the $\mu_{a_n}(\lambda)$ 5 vector output from equation 8.30 and the resulting $\mu_a(\lambda)$ 5 vector from the tissue absorption due to the modelled absorbers (Hb, HbO_2, H_2O, and Lipids fitted using least squares). The polynomial de-trending meant that all raw intensity data were normalized to a reference level of 1. This is also the reference intensity in equation 8.30, so as $ln(1/1) = 0$ it can be seen that the fit error fraction would be

expected to drop to around zero in the final five seconds of each ten second period, as it was over these final five seconds that both the polynomial de-trending and the pulsatile modulation depth fitting were conducted, so the solution $\mu_a = \mu_{a_{ref}}$ gives a least squares fit solution with no errors.

It can be seen from the baseline tissue molarities in figure 8.11 that the baseline tissue saturation is very low, around 16%. This seems unreasonable given that most published studies have found baseline tissue saturations in the range of 60% to 80% (for example data refer to Hyttel[225]). There is no significant change in tissue Hb during the refill period, which also seems infeasible, especially considering the significant amount of Hb present. The tissue H_2O does not significantly change either. It is unclear to what degree interstitial fluid flow would be expected, so the lack of significant change in water volume fraction is not necessarily unexpected. The baseline fractional volume of water is well within the typical range.

It can be seen that if the Hb and HbO_2 molarities are added to obtain a total molarity Hb_{total}, the rate of change of total haemoglobin will be equal to the rate of change of HbO_2. From figure 8.11 it can be seen that immediately after the pressure release, this gradient is approximately $2 \times 10^{-5} Moles L^{-1} s^{-1}$. Assuming a total blood haemoglobin molarity of 2.3mMolar, this is equivalent to a perfusion volume fraction of 8.7×10^{-3} per second. This corresponds to a perfusion of approximately 52ml/100g/minute, around an order of magnitude higher than published values for the healthy adult forearm (for example as found by Blasi et al [165]). This suggests that either the inversion technique is grossly inaccurate, or that capillary refill does not measure the perfusion of blood.

A recording of multiple capillary refills was made on a second volunteer, and processed in identical fashion to yield changes in tissue molarities as a function of time.

It can be seen that (figure 8.12) although the fit errors are roughly the same between the two datasets, in the second volunteer the HbO_2 drops to negative values in some refills. This is clearly not feasible. The baseline tissue saturation is extremely high at around 98%, which is above the expected range. Comparing total baseline haemoglobin between the two datasets there is a very large difference, with the first set having a molarity of approximately 6×10^{-4} and the second 1×10^{-5}. This corresponds to blood volume fractions of 26% and 0.44% respectively.

Brix et al [226] used a CT scanning technique in combination with an injected bolus of tracing agent to determine blood volume fraction in adult neck muscle as 4±1%, so both of these volume fractions would appear to lie outside the normal range.

When new data were recorded from the same site on the Thenar Eminence of the first volunteer, it was found that the modulation depth fitter converged to a solution with

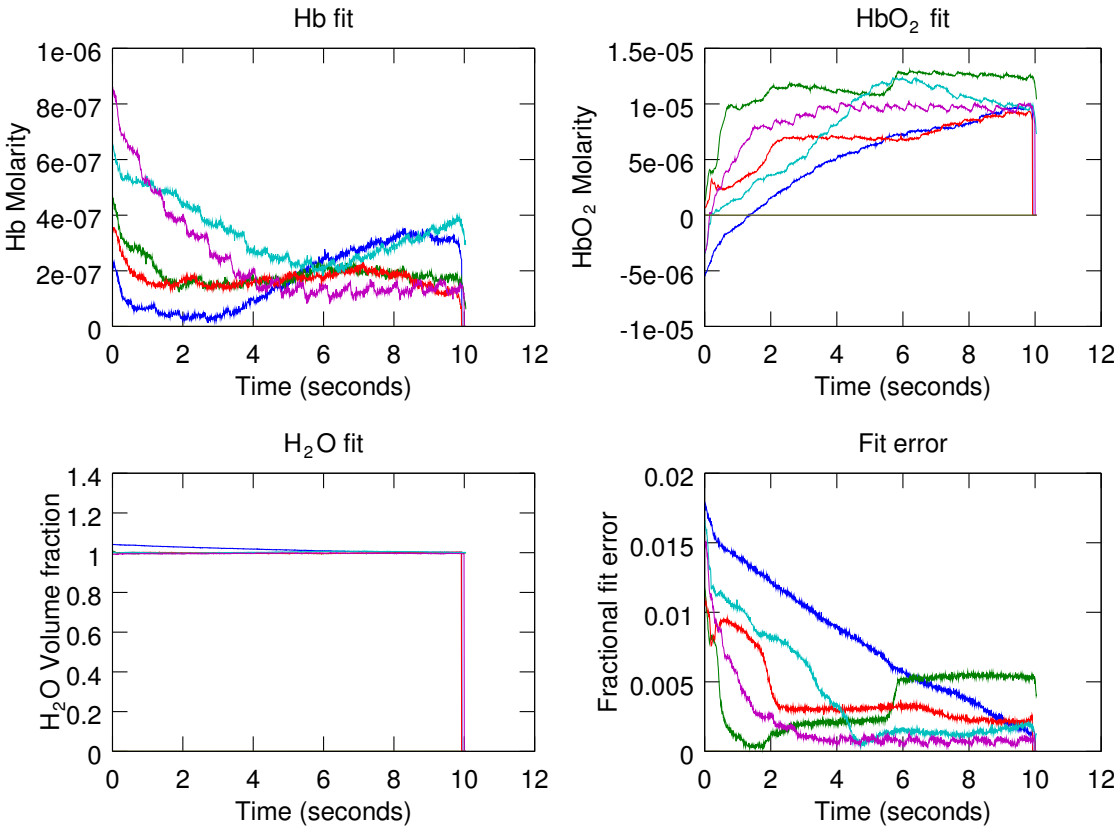


Figure 8.12: NIRS capillary refill data recorded from a second volunteer, each line is a single refill.

negative oxy-haemoglobin. Limits were added to the sqp optimiser function to prevent this occurring, but the optimiser was still found to converge to the limiting value of a zero oxy-haemoglobin molarity. Recording a dataset consisting of multiple recordings on subsequent days from the first volunteer under identical conditions, it was found that the modulation depth fitting was both highly unreliable (failing to reach a solution without reaching a limit in one of the fitting dimensions in 60% of cases), and highly variable from recording to recording when it did converge, with both the blood volume fraction and tissue saturation changing apparently randomly over the entire range of possible solutions.

The reasons for the poor performance of the modulation depth fitting technique are unclear. The reflection mode mono-chromator results yielded relatively consistent results that were not infeasible when processed with the same fitting algorithm (refer to Chapter 6). However, when applied to the five wavelength NIRS data, the model was no longer highly overdetermined, and this may be one cause of the poor stability. A second and potentially more significant difference between the two datasets was the sensor design. Although both devices used reflection mode NIRS with a roughly equivalent source to detector separation, the mono-chromator used a single fibre bundle for illumination,

whereas in the LED system, each wavelength follows a different path length through the tissue. Thus in the LED based system any spatial inhomogeneity in the pulsatile volume fraction (ΔV in equation 6.20) will make the fitting model inaccurate. Rubins et al found significant inhomgeneity in the pulsatile modulation depth measured using diffuse reflectance spectroscopy on the palm of the hand [227]. Although the diffuse reflectance path length is much shorter than with the custom built NIRS sensor used for the LED studies (as illustrated by figure 6.3), it can be seen that the significant ($\pm 50\%$) changes seen by Rubins et al would be great enough to cause poor fitting performance.

In conclusion it would seem that the effect of the different optical paths on pulsatile modulation depths in conjunction with the poor tolerance of the fitting algorithm to errors leads to the inconsistent results with this algorithm. These problems are further exacerbated as the number of degrees of freedom in the data are only slightly greater than those in the model, meaning that small errors in modulation depth at a single wavelength are sufficient to cause significant fitting errors.

For these reasons it was determined that the modulation depth technique was unsuitable for use with the NIRS sensor, although an improved sensor with closely spaced LEDs (bare die construction or chip on board is used in pulse oximeter probes) may prove successful.

8.2.2 Fitting a diffusion model

The diffusion model was described earlier in Chapter 6 in the context of modelling the pulsatile modulation depth. The previous analysis leading to the derivation of equation 6.20. However, if an appropriate set of assumptions are made (for example that transport of only haemoglobin and oxy-haemoglobin out of and into the pressure affected area are observed) then the same model can be extended and fitted to the capillary refill directly.

From equations 8.26 to 8.29:

$$ln(I) = ln(I_0) - d\sqrt{3\mu_a(\mu_a + \mu'_s)}. \tag{8.31}$$

If a change of unit system is made such that d is redefined as being equal to 1, the equation can then be reformulated in this new unit system, defining $ln(I_0) = B$, where B is a constant representing the source intensity:

$$ln(I) = B - \sqrt{3\mu_a(\mu_a + \mu'_s)}. \tag{8.32}$$

Making the assumption that $\mu'_s \gg \mu_a$:

$$ln(I) = B - \sqrt{3\mu_a\mu_s'}. \tag{8.33}$$

Moving to a multiple wavelength vector space:

$$\vec{\mu_a} = \frac{(\vec{B} - ln(\vec{I}))^2}{3\vec{\mu_s'}}. \tag{8.34}$$

This equation can be applied to the raw intensity data, with the set of coefficients $B(\lambda)$ (or $\vec{B}$) being optimised to minimise the least mean squares fit error to a set of absorbers. The intensity values in raw units can be used in this equation, as $ln(aI) = ln(a) + ln(I)$, meaning that optimisation of the coefficient B can be used to account for the unknown detector sensitivity as well as the unknown reference intensity.

Expanding equation 8.34:

$$\vec{\mu_a} = \frac{\vec{B}^2 - ln(\vec{I})^2 - 2\vec{B}ln(\vec{I})}{3\vec{\mu_s'}}. \tag{8.35}$$

If a polynomial re-normalization technique is applied to scale the raw intensity data to a reference level of 1, then as $ln(1) = 0$, the $ln(\vec{I})$ values will be centred around zero. In tests of the capillary refill device on healthy volunteers, the maximum deviation in intensity was typically $\pm 20\%$. As this implies that $ln(I_{final}/I_{initial}) = ln(1.2) = 0.18$, it can be seen that the $ln(\vec{I})^2$ term will be small compared to the $ln(\vec{I})$ term, and so equation 8.35 can be simplified to:

$$\vec{\mu_a} = \frac{\vec{B}^2 - 2\vec{B}ln(\vec{I})}{3\vec{\mu_s'}}. \tag{8.36}$$

If the assumption is made that the baseline tissue absorbers are unknown, then the values of the $B(\lambda)$ coefficients are essentially arbitrary unless the absolute intensity of the sources is calibrated. However, if the model is instead used in a differential form, then it is possible to reformulate the expression in a way that may allow B to be found. If a baseline $\vec{\mu_a}$ is defined at the reference level where $ln(\vec{I}) = 0$, then:

$$\Delta\vec{\mu_a} = \frac{-2\vec{B}ln(\vec{I})}{3\vec{\mu_s'}}. \tag{8.37}$$

It is not feasible to solve this equation for $B(\lambda)$ under all conditions. For example, if the observed changes in $\vec{\mu_a}$ are due to a changing concentration of a single absorber, or to the concentrations of multiple absorbers changing in a constant ratio, then a $B(\lambda)$ 5 vector can always be found to match the observed changes in $ln(\vec{I})$ to an absorber from the modelled set of absorbers. An even more fundamental problem is that there

is no means to estimate the magnitude of the $B(\lambda)$ 5 vector, and therefore no means to obtain absolute concentrations of absorber from this model. It should also be noted from equation 8.37 that even if it proves possible to determine the $B(\lambda)$ coefficients, there is still a dependence on the absolute value of $\vec{\mu_s'}$, so if absolute changes in Hb and HbO_2 molarities are desired, $\vec{\mu_s'}$ must either be estimated or found by another independent method. However, if only the relative changes in Hb and HbO_2 are needed, then a power law approximation for $\vec{\mu_s'}$ combined with an $B(\lambda)$ fit constrained to a fixed vector length would be suitable.

If there are multiple degrees of freedom within the input data, if for example two absorbers are changing in different ratios throughout the dataset, then it may be possible to reach a solution in arbitrary concentration units using this method. For the $\Delta\mu_a(\lambda)$ measured absorptions from equation 8.37 to be mapped to a set of absorbers, the number of absorbers must be equal to or smaller than the number of wavelengths. If the number of absorbers is less than the number of wavelengths (i.e. 5 in the case of the custom made NIRS sensor), then it can be seen that if the number of wavelengths is n, then the range of possible absorptions from the modelled absorbers forms a sub-space with $n-1$ or fewer dimensions within the n dimensional space of measured absorbency $\mu_a(\lambda)$. Optimisation of the n space vector $B(\lambda)$ minimizes the mean squared distance between the $\mu_a(\lambda)$ data-points and this subspace.

In the case of two of more absorbers always varying together in a constant ratio, then it can also be seen that the number of dimensions of the sub-space is reduced. Thus by extending this dimensionality argument it can be seen that the number of degrees of freedom to identify absorbers is equal to or less than the number of degrees of freedom in the data. This is not the same as the number of wavelengths, for example if $ln(I(\lambda))$ for a three wavelength system was found to produce a coplanar point cloud, then there would only be two degrees of freedom within the data.

Application of equation 8.37 to the datasets recorded from the automated capillary refill device was attempted. The sqp optimisation function in GNU-Octave was used to optimise the $B(\lambda)$ vector such that the resulting $\vec{\mu_a}$ could be fitted to a change in Hb and HbO_2 molarity using a least squared fit with minimal error. This was found to produce trivial solutions, as the changes in $ln(I(\lambda))$ mostly fell along a straight line in five λ space. This resulted in the $B(\lambda)$ vector settling on the first value that scaled the $ln(I(\lambda))$ vector onto an absorber, i.e. a trivial solution that bore no relation to the actual absorbers that were changing in concentration during the refill.

Returning to equation 8.36, it can be seen that if both the baseline absorbers (and thus $\mu_a(\lambda)$) and the transport scattering coefficient for the tissue $\vec{\mu_s'}$ are known, then $\vec{B}$ can

be found directly:

$$\vec{B} = \sqrt{3\vec{\mu_s'}\vec{\mu_{abaseline}}}. \tag{8.38}$$

Substituting back into equation 8.36:

$$\vec{\mu_a} = \frac{3\vec{\mu_s'}\vec{\mu_{abaseline}} - 2\sqrt{3\vec{\mu_s'}\vec{\mu_{abaseline}}}ln(\vec{I})}{3\vec{\mu_s'}}, \tag{8.39}$$

moving back to a differential form of the equation:

$$\Delta\vec{\mu_a} = \vec{\mu_a} - \vec{\mu_{abaseline}} = \frac{-2\sqrt{3\vec{\mu_s'}\vec{\mu_{abaseline}}}ln(\vec{I})}{3\vec{\mu_s'}}, \tag{8.40}$$

$$\Delta\vec{\mu_a} = \frac{-2\sqrt{\vec{\mu_{abaseline}}}ln(\vec{I})}{\sqrt{3\vec{\mu_s'}}}. \tag{8.41}$$

Here the $\vec{\mu_{abaseline}}$ arises from a combination of the baseline blood (Hb + HbO_2) and the tissue absorption. The custom made NIRS sensor does not use wavelengths at the main lipid absorption peak at approximately 920nm, but the 950nm LED is likely to encounter significant attenuation from tissue water. Assuming 100% water volume fraction in tissue, and a 20mm effective path length through the tissue at 950nm, then there will be around 20% absorption along the path. Thus it can be seen that a model for $\mu_{abaseline}(\lambda)$ including a fixed volume fraction of water (90%) might be appropriate for use on areas without significant subcutaneous fat. There would also be significant baseline absorption from both haemoglobin and oxy-haemoglobin, the molarities of which would be highly variable between measurement sites and experiment volunteers. Thus it can be seen that there are two degrees of freedom in this new fitting problem; the Hb and HbO_2 molarities. These can be optimised to find the values that result in the minimum error when $\Delta\mu_a(\lambda)$ is fitted to a change in Hb and HbO_2 molarity using a separate least squares fitter, a repeat of the process used previously to find the $B(\lambda)$ vector.

This new approach reduces the number of degrees of freedom in the fitting process from 5 to 2, and should reduce some of the tendency to settle on trivial solutions. More importantly, the fixed volume fraction of water acts as a constraint on the magnitude of the $\mu_{abaseline}(\lambda)$ vector, resolving the arbitrary units problem encountered earlier. It can be seen that as the concentration of water is constant, an increase or decrease in the total haemoglobin molarity will distort the absorption spectrum, meaning that it may be possible to analyse total haemoglobin Hb_{total} using this method, provided that

$\vec{\mu_s'}$ can be estimated. In this respect it is also helpful that the $\vec{\mu_s'}^{-1}$ dependence is now a $\vec{\mu_s'}^{-1/2}$ dependence, approximately halving the impact of errors in its estimation.

In the analysis so far, the impact of the de-trending operation on the raw data has not been thoroughly analysed. If the baseline drift is principally due to variations in the LED output power or the detector sensitivity, then it can be seen that the use of de-trending is valid, as the differential model of $\Delta\vec{\mu}_a$ assumes that there is a consistent baseline intensity level I_0, and that all intensity deviations arise due to changing absorber concentrations. The previous use of de-trending (in the pulsatile modulation depth precessing) was based on the assumption that all long term drift was due to intensity and or sensitivity fluctuation, and there was a consistent baseline tissue attenuation, which could be derived from the modulation depth fitting algorithm applied over the entire recording duration (in windows away from pressure applications). The short term drift due to changes in absorber concentration was assumed to be confined to individual capillary refills, with each refill returning to the same baseline level of tissue attenuation.

If however the tissue baseline absorption drifts over the course of the data recording, for example due to changes in sensor position, or due to regulation of the vasodilation due to the sympathetic nervous system, then the use of de-trending can be seen to be inappropriate if these effects dominate over changes in illumination intensity. If there are significant changes in both baseline absorption and LED intensity over the course of the recording, then it not possible to obtain useful data unless an independent measure of the changing illumination intensity is provided.

Considering the performance of the front-end amplifier (the dual op-amp design was used with the F4-Discovery NIRS system), it would appear that sensitivity drift will be insignificant. The amplifier gain is set by the feedback impedance and the bandpass gain, and with the likely temperature drift over the course of the experiment in the 1°C range (as the daughterboard is separated by around 2cm from the F4-Discovery and consumes negligible current, meaning temperature shifts will likely only be those of the room air), the drift in both of these gains will be considerably less than 0.5%, as demonstrated by the earlier analysis of the effect of front-end component tolerance.

However, changes in LED quantum efficiency due to changing temperature and LED drive voltage may be significant. In order to investigate the likely magnitude of such drift effects, the NIRS device was operated for 45 seconds, starting from "cold" with a block of highly scattering polystyrene foam placed underneath the sensor.

From figure 8.14 it can be seen that there is a small warm up transient over the first 10 seconds, corresponding to a change of approximately 0.4% maximum intensity in all the NIRS channels except 810nm, where the effect is much smaller. The transient

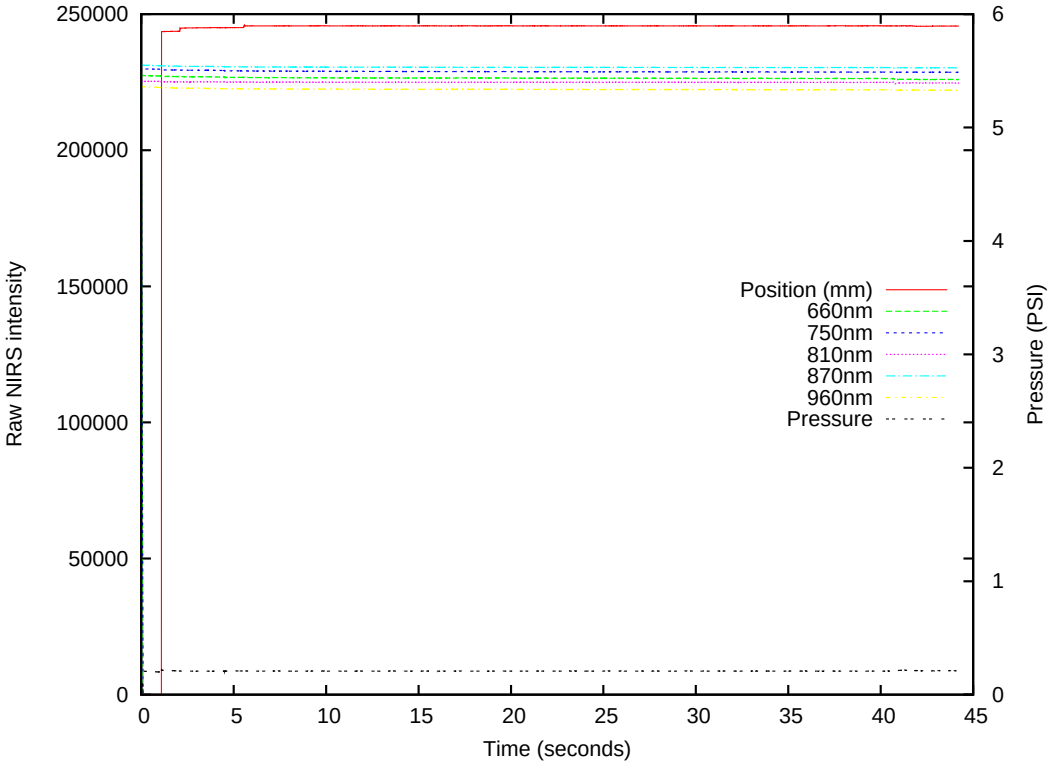


Figure 8.13: NIRS data recorded from a foam block using the automated pressure device.

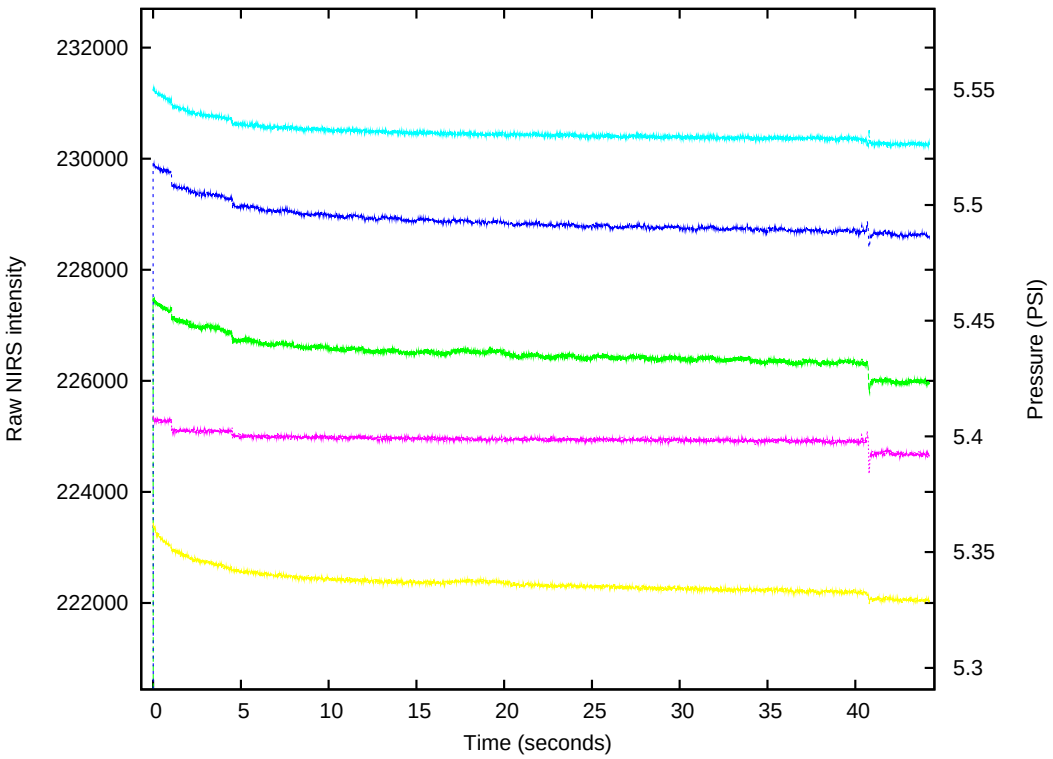


Figure 8.14: Detail of the NIRS intensity data from figure 8.13.

drops off relatively rapidly with a change of order 0.2% in intensity over the next 35 seconds. Thus with reference to the pressure induced tissue absorption changes seen in figure 8.10, it can be seen that changes in LED intensity and detector sensitivity are insignificant compared to the tissue changes, especially after the first few tens of seconds after the device is turned on. The signal to noise ratio is also evident from figure 8.14, and is very close to 14 bits per sample, only slightly lower than the 15 ENOB/sample derived in equation 7.27. The data in figure 8.14 were recorded using a late revision "Z" STM32F4, it was found that earlier silicon revisions suffered from significant ADC noise with the instruction cache enabled, reducing the ENOB unless the cache was turned off.

It should be noted that if the device is used on a volunteer instead of a foam block, the skin surface temperature is likely to be significantly warmer than room temperature, and so the temperature of both the LEDs and photodiode will increase. As the thermal diffusivity of silicone rubber is approximately $10^{-7} m^2 s^{-1}$, and the LEDs are located under approximately 1mm of silicone rubber ($10^{-3} m$), the time constant for LED warmup can be estimated. From the Fick equation for diffusion:

$$t_D = \frac{L^2}{4D} = \frac{10^{-6}}{4 \times 10^{-7}} = 2.5\, s. \tag{8.42}$$

Thus it can be seen that the silicone encapsulation is sufficiently thin that thermal equilibrium with the skin surface will be reached over a timescale of order 10 seconds, and thus the heating of the LEDs via the skin will only cause drift at the start of each recording session.

Refractive index mismatching due to air gaps between the sensor surface and the skin is another source of drift in intensity. Equation 7.28 shows that there is a potential reflection loss of more than 20% of the incident light at each transition, comparable in magnitude to the observed refill signals. As this intensity loss takes place at the interface between the LEDs and photodiode, it can be seen that it is indistinguishable from changes in source intensity or detector sensitivity, and thus re-normalization of the data would be an appropriate method of combating this effect. However, as coupling changes are likely to be mainly a function of applied pressure rather than time, especially if the volunteer's arm or hand is held steady under the sensor throughout the recording, it can be seen that coupling induced changes are likely to be most significant over the source of each pressure application and release cycle. Therefore application of any form of de-trending is impossible, and the only method that may be used to reduce coupling changes is application of an optical coupling solution to the skin surface.

In figure 8.10, the sensor can be seen to have momentarily lost contact with the skin surface around 70 seconds into the recording (visible as a large position jump). This

causes a significant decrease in intensity due to refractive index mismatch. The light level returns to the original intensity once the sensor adjusts to the new hand position, but the optical intensities have changed slightly. These changes may well be due to the position having moved slightly to a region of tissue with differing haemoglobin content.

Thus in conclusion it would appear that use of polynomial de-trending is inappropriate for application of the diffusion model described by equation 8.41 to the automated capillary refill NIRS data. From consideration of the data in figure 8.10 and measurement of the drift in LED intensity after the device is turned on, it would appear a suitable data processing route would be exclusion of the first 15 seconds of data, followed by simple normalization of each optical channel by dividing by its mean value. Throughout the recording period, as much consideration as possible should be put into maintaining contact between the sensor and the skin surface, so as to avoid coupling artefacts. A coupling solution or gel, such as glycerol should be used to minimize the potential for coupling problems to arise.

The normalized intensity data would then be in a suitable form for the fitting algorithm based on equation 8.41 to be applied.

In order to ascertain whether application of this model was indeed feasible, a script was written in the GNU-Octave language to implement optimisation of the baseline tissue Hb and HbO_2 molarities using the sqp optimiser. Minimization of the error between a least squares fit of Hb and HbO_2 absorption vectors to $\Delta\vec{\mu}_a$ could then be attempted through optimisation of the baseline Hb and HbO_2 with volumetric water held constant (as described earlier). As with the earlier data processing algorithm which used pulsatile modulation depth, published time and distance resolved studies were used to estimate a typical absolute value for μ_s' for use in 8.41, using equation 6.17 with a k value such that μ_s' was equal to $0.45mm^{-1}$ at 800nm wavelength. Equation 8.41 is only applicable in a unit system with unit length equal to d, the source to detector separation. Converting back into S.I. units:

$$\Delta\vec{\mu}_a = \frac{-2\sqrt{\vec{\mu}_{a\,baseline}}ln(\vec{I})}{d\sqrt{3\vec{\mu_s'}}}. \tag{8.43}$$

This equation is the same as that derived by Chance et al [228]. Boas [229] derives an expanded form of this equation that does not make as many assumptions regarding the relative sizes of the coefficients:

$$\Delta\vec{\mu}_a = \frac{-2\sqrt{\vec{\mu}_{a\,baseline}}}{d\sqrt{3\vec{\mu_s'}}}\left[1 - \frac{1}{1 + d\sqrt{3\vec{\mu}_a\vec{\mu_s'}}}\right] ln(\vec{I}). \tag{8.44}$$

It can be seen that this expanded form in equation 8.44 approaches the simplified from of 8.43 as the total separation from source to detector (d) exceeds several optical depths of tissue.

A recording was taken from the palm of a healthy volunteer's hand during repeated 3PSI pressure applications over a 15 second period (5 seconds of 3PSI applied, followed by 10 seconds of 0.2PSI, normalized raw intensity shown in figure 8.15).

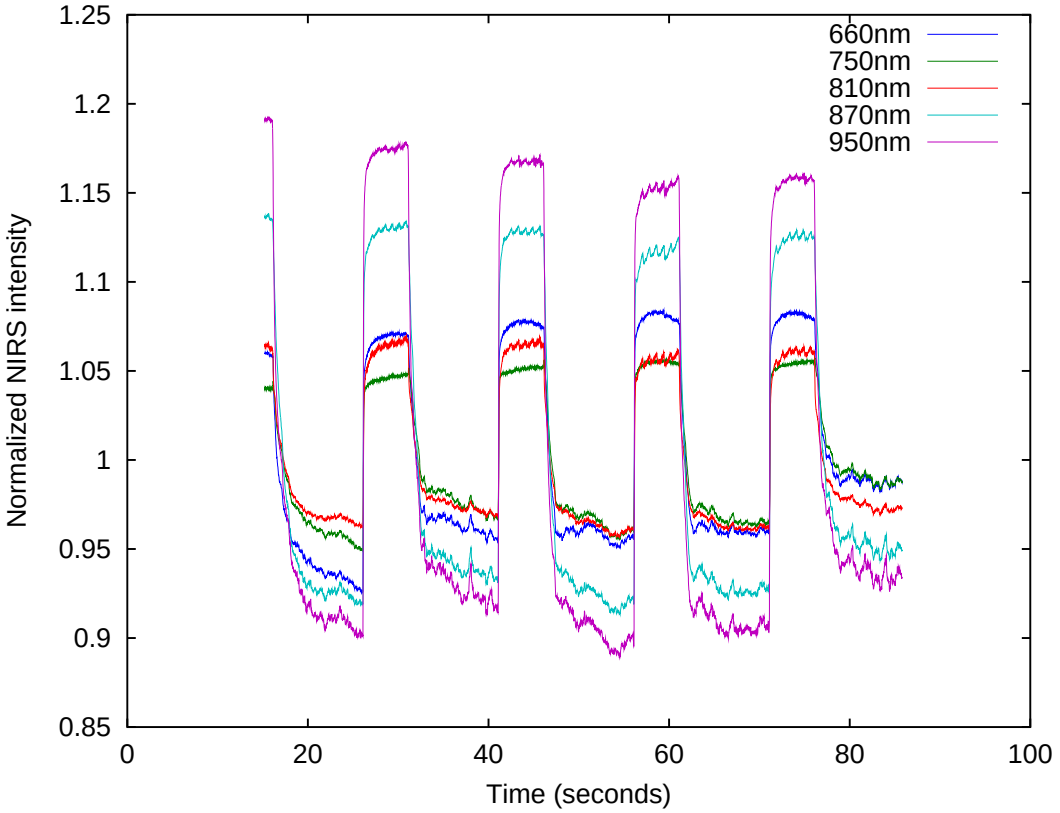


Figure 8.15: Sample refill data used to benchmark model fitting methods (using second healthy adult volunteer).

The raw data were first cropped and normalized, using the procedure described previously, and was then used to test the performance of various implementations of fitting to equation 8.43. As the majority of the data are clustered around two values - the baseline values during pressure application and no pressure application, it was determined that use of the mean fractional fit error may not be an optimal error metric. For this reason the 95% quantile level of the fractional μ_a error over the dataset was used as a goodness of fit measure, as this should form a measure that is most strongly weighted towards the worst case fit error. Defining a 5 by 2 component absorption matrix $\mathbf{A}$ and a 2 component vector of Hb and HbO_2 molarity changes m:

$$\Delta\vec{\mu}_a = \mathbf{A}\vec{m}. \tag{8.45}$$

The matrix A was generated from published haemoglobin and oxy-haemoglobin absorption spectra combined with the measured LED spectra for each LED channel (as shown in figure 7.12), using a numerical integration proceedure as outlined previously. The ordinary least squares solver in the octave environment was used to solve this equation for each candidate. In this relatively simple case, ordinary least squares is equivalent multiplication of the $\Delta\vec{\mu}_a$ vector with the Moore-Penrose pseudoinverse of the matrix $\mathbf{A}$, i.e.

$$\vec{m_{fit}} = \mathbf{A}^\dagger\Delta\vec{\mu}_a, \tag{8.46}$$

where $\mathbf{A}^\dagger$ is defined as $\mathbf{A}^\dagger = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T$. As the estimation change in absorption coefficient is a five vector and there are only two changing absorber concentrations in this model, it is possible to define an error vector $\vec{E}$, as an exact fit will not be possible:

$$\vec{E} = \Delta\vec{\mu}_a - \mathbf{A}m_{fit} = \Delta\vec{\mu}_a - \mathbf{A}\mathbf{A}^\dagger\Delta\vec{\mu}_a, \tag{8.47}$$

$$\vec{E} = \Delta\vec{\mu}_a - \mathbf{A}(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\Delta\vec{\mu}_a, \tag{8.48}$$

$$\vec{E} = \left[\mathbf{I} - \mathbf{A}(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\right]\Delta\vec{\mu}_a. \tag{8.49}$$

Introducing the sample index n, a fractional error for each sample point was then defined as:

$$f_n = \frac{|\vec{E_n}|}{|\Delta\vec{\mu}_{a_n}|}. \tag{8.50}$$

The goodness of fit metric for the entire dataset was then defined using the quantile Q as:

$$fitness = Q_{0.95}(f), \tag{8.51}$$

where Q_τ is the level τ quantile.

Finding the maximum fractional fit error returns the absolute worst case scenario, but use of the maximum would be prone to being strongly influenced by any momentary artefacts in the data, so the 0.95 level was chosen. This error metric was calculated

as a function of tissue saturation and (log scale) tissue total haemoglobin used in the calculation of $\vec{\mu_{abaseline}}$ together with an assumed 90% water content. Tissue molarity is typically in the region of 10^{-4}. However, this varies greatly from person to person, and for this reason a logarithmic scale covering three orders of magnitude was used in the plots.

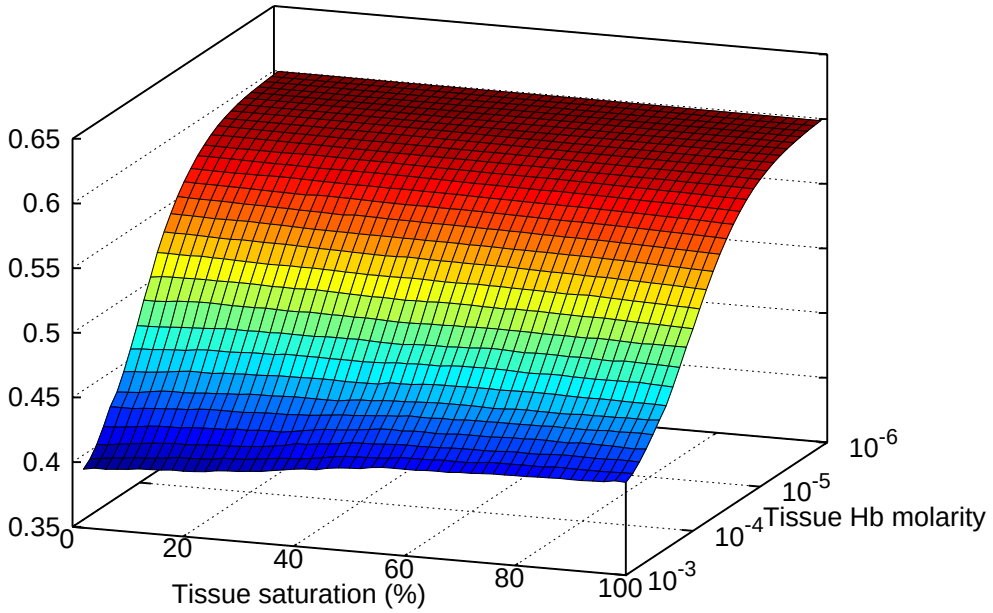


Figure 8.16: Fit error metric using the 95% quantile level. Hb molarity is total haemoglobin molarity.

As can be seen from figure 8.16, the best fit value is at a saturation of 0% and an unreasonably high tissue molarity at the limit of the graph. One possible explanation for this poor result is that application of pressure causes significant structural changes in the tissue. For this reason the goodness of fit function was plotted a second time, this time with all the pressure application regions excluded from the data, and with the new set of data which was recorded with a constant applied pressure of 0.2PSI normalized using the same procedure as previously.

It can be seen from figure 8.17 that this procedure produced very similar results, although there is some more apparent noise in the fit function, and the fractional error is considerably higher. The reason that the best fit to the data is at an unrealistic tissue molarity and saturation is unclear. One possible problem can be seen from consideration of equation 8.50, the vector magnitude of the absorption coefficient five vector and error

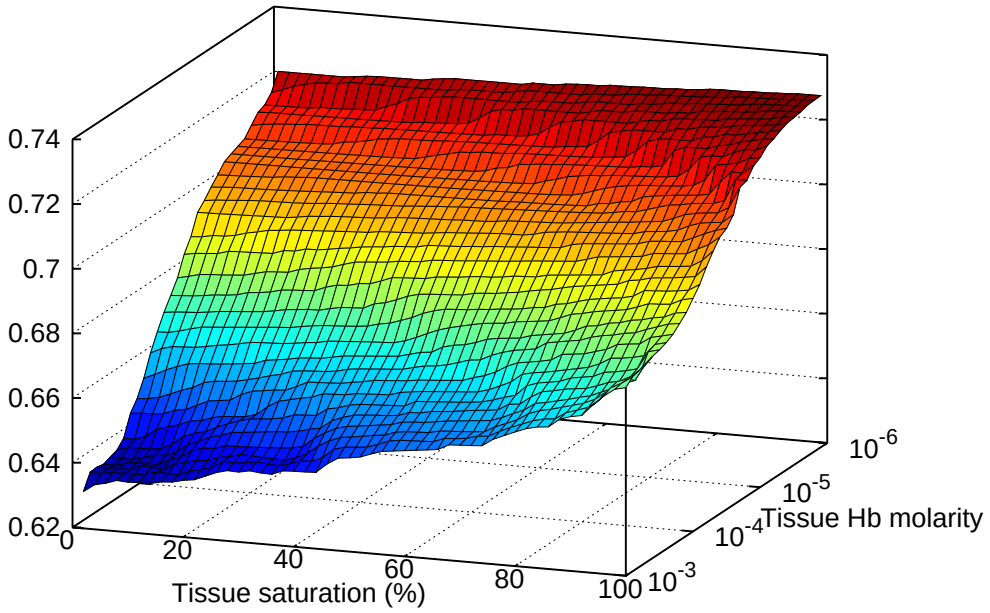


Figure 8.17: Fit error metric using NIRS data from 0.2PSI applied pressure time periods only. Hb molarity is total haemoglobin molarity.

five vector are used. Thus if there is a large fractional error in the fit to one of the $\Delta\vec{\mu}_{a_n}$ components, but the magnitude of this $\Delta\vec{\mu}_{a_n}$ component is small compared to the other vector components, the overall contribution to the error metric will be small.

For these reasons plot 8.17 was repeated with equation 8.50 replaced by an element wise fractional error calculation. The set of fit errors f_n used in equation 8.51 was now a larger set of fractional error values; one for each intensity data point, rather than one for each multi-wavelength intensity vector sample point. The resulting error is shown in figure 8.18. It was noted that at intensities close to the mean level, where ln of the intensity was close to zero, the effect of artefacts and noise was greatly increased. For this reason data points were excluded when the change in intensity at any of the wavelengths was less than 1%.

From the figure, it can be seen that there is less significant rippling in the plot surface as a result. The best fit value now occurs at a tissue saturation of 84%, a credible saturation. However, the fit still has an excessively high total tissue molarity. The reason for this would appear to be the high modulation depth of the refill signal at 950nm. The large intensity change means a relatively large path length is required,

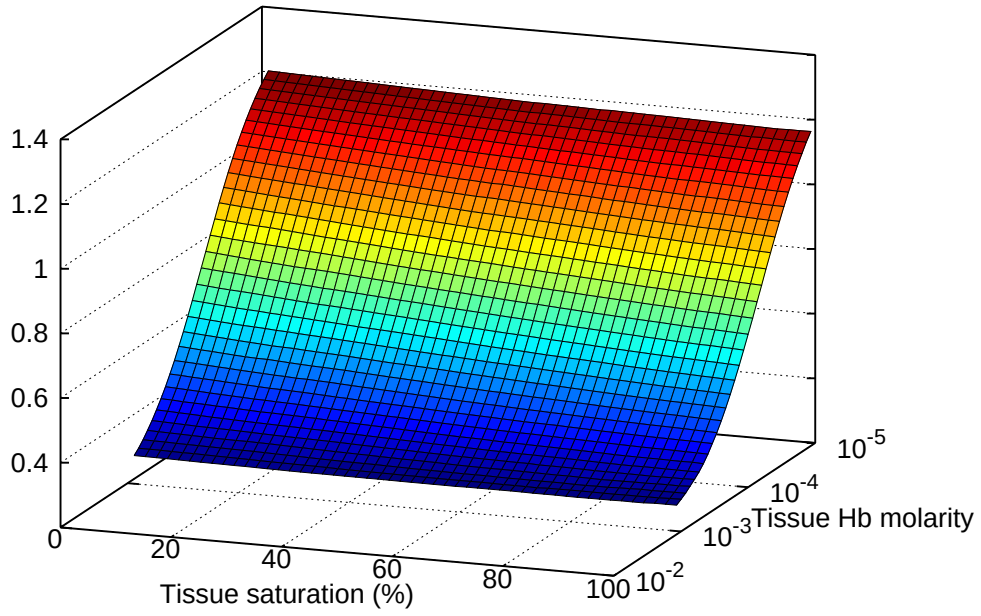


Figure 8.18: Fit error metric using NIRS data from 0.2PSI applied pressure time periods only. Hb molarity is total haemoglobin molarity. Repeat of the previous plot with error metric based off element wise fractional errors.

meaning a better fit is obtained when the absorption from water is insignificant. By extending the range of molarities to around 2 millimolar, it was found that the best fit fractional error asymptotically approached a value of approximately 0.21 at 84% tissue saturation, confirming the hypothesis that inclusion of water decreases the quality of the fit to the data. The reason for this is unclear, but it could be a simple case of the tissue inhomogeneity combined with the placement of LEDs. An alternative suggestion would be that the presence of other absorbers within the tissue (such as cytochrome oxidase) leads to a more uniform "background" absorption, and that the presence of water in the model decreases accuracy as water only absorbs significantly in the 950nm wavelength channel.

In order to investigate the validity of these hypothesis further, the wavelength dependent path length factor as measured by Essenpreis et al [149] and the BORL group at UCL was compared to the model in equation 8.43. This path length factor (sometimes refereed to by the acronym DPF for Differential Path Factor in the literature) is the expectation value for the optical path length through tissue, which can be seen to be equal to:

$$DPF = \frac{ln(\vec{I})}{d\Delta\vec{\mu_a}} = \sqrt{\frac{3\vec{\mu_s'}}{4\vec{\mu_a}_{baseline}}}. \tag{8.52}$$

In order to compare this equation with the published studies, a tissue total haemoglobin molarity of 115 micromolar and 84% saturation was first assumed, in line with published data.

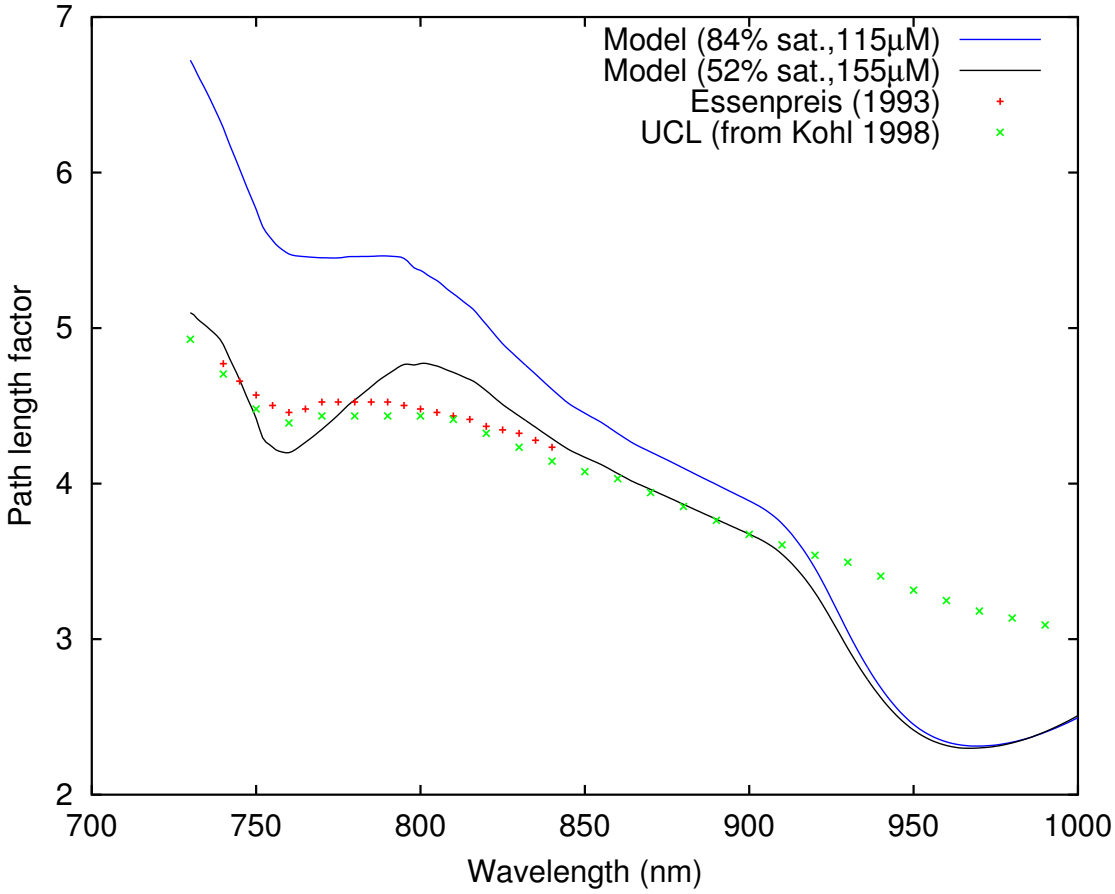


Figure 8.19: Path length factor computed from equation 8.43 and typical tissue haemoglobin molarity and saturation.

It can be seen that the theoretical path length in figure 8.19 (blue line) is a relatively good match to the shape of the Essenpreis and UCL data, although the decrease around 750nm is not as pronounced in the Essenpreis study. The data sourced from the BORL group at UCL and presented by Kohl et al in [148] covers a greater wavelength range and is plotted as the green stars. It would seem that the overall magnitude of the path factor from the model (using a power law with coefficient -0.65 and an absolute value of 4.5/cm for μ_s') is very similar to that found in the time resolved studies, but it is also clear that the significant dip in path length factor at 950nm is not present in the UCL data, and that the overall gradient of the decrease in path length with wavelength is considerably higher in this initial modelling attempt. For these reasons a least squares

regression was used to fit the model to the UCL dataset by adjusting the tissue total haemoglobin and saturation. This resulted in a considerably lower saturation of 52%, somewhat below the lowest values seen in healthy tissue. The total haemoglobin molarity increased to 175 micromolar, which is somewhat higher than typical. This new model (black line) would seem to be a closer fit to the published path length factors, with the overall downward trend with increasing wavelength presented in figure 5 of the Kohl et al publication [148] matching the black line well. However, there is still a significant drop in the modelled path length around 950nm due to the water absorption peak, and the high level of haemoglobin in the model fit causes a significantly more pronounced dip in path length around 750nm than is present in the published data.

The lack of a pronounced dip in the path length factor at 950nm explains the poor performance of the model fitting attempted in figure 8.18, as a water volume fraction of approximately zero in the model would remove the path length decrease around 950nm, giving a better fit to the data. However, a high percentage of water is clearly present in tissue, so this is not an adequate explanation of the observed effects. The likely explanation would appear to be the presence of other absorbers within the tissue, for example cytochrome oxidase is known to be a relatively significant absorber in the near infra-red. Chance et al [230] determined the near infra-red absorption spectra for the various cytochrome species. From figures 1 and 2 in their publication, it can be seen that cytochrome absorption increases with decreasing wavelength, from a minima at around 950nm. Thus it would appear that absorption by cytochrome oxidase can explain the observed path length factor, with a tissue saturation in the region of 80% and a total tissue haemoglobin concentration in the region of 100μM (values similar to those reported for the healthy adult forearm), the path length will be similar to that plotted with the blue line in figure 8.19. Inclusion of cytochrome oxidase species into the model will result in a decrease in the path length at lower wavelengths, giving a theoretical path length that much more closely matches the published data.

Unfortunately inclusion of other absorbers increases the number of degrees of freedom in the model fitting process, making fitting to the 5 wavelength data from the custom sensor impossible. However, it would appear that use of water as an absorption reference to quantify absolute concentration of other absorbers is a feasible technique, for example Matcher et al described a similar method in a 1994 paper [231], although a slightly different method based on identification of spectral features as multiple wavelengths was used in this case.

As the application of equation 8.43 requires a more complex model of tissue absorption than the data will allow, it was determined that application of a more simple differential

path model based on published data for applicable tissue types may be more suitable for the five wavelength NIRS data. To this end, the differential path length model was researched further, with a particular emphasis on quantification of errors.

8.2.3 Fitting a differential path length model

The differential path length model was introduced previously in Chapter 6 (covering optical methods). This is a method which would seem especially well suited to study of the automated capillary refill data, as it requires multi-wavelength intensity data recorded at only a single source to detector separation. The fact that only a differential measure of absorber concentration can be obtained is not a significant problem for use in capillary refill studies, as the effect of the applied pressure change on haemoglobin and oxy-haemoglobin is the most scientifically useful question to be answered. Specifically, the rate of change of total tissue haemoglobin molarity following the release of pressure can be easily obtained via the differential path length method, meaning that the rate of change in blood volume fraction in the tissue can be found (if assumptions are made regarding blood haemoglobin molarity). This is the most useful metric for use in determining the nature of the capillary refill, as if the refill is due to inflowing arterial blood, then this volumetric perfusion should match that derived by other techniques, such as venous occlusion.

As previously outlined, the path length factor is the ratio of the experimentally derived path and the source to detector separation. The differential path length technique relies on use of an experimentally derived wavelength dependent path factor estimate, usually obtained via time resolved hardware using pulsed laser illumination combined with a photo-multiplier tube as a sensitive and high bandwidth photo-detector. Thus unless time resolved hardware is used to find the path length at the measurement site on each volunteer, a technically challenging solution, it is necessary to use an estimated path length factor based on published tissue data. Usually a compromise is made, and published path length data that are specific to the tissue type under consideration is used under the assumption that inter subject variability will be relatively low.

Duncan et al [147] found path length factors in the forearm using a time resolved modulated laser array. In a population of 100 healthy adult volunteers (50 male and 50 female), a standard deviation of approximately $\pm 20\%$ of the path length was found across the entire group of 100 subjects. This decreased to approximately $\pm 15\%$ standard deviation if male and female subjects were separated. Similar levels of variance were found in the adult leg when measurements were made of the calf tissue, and there was only a small change in the variance between the four different laser wavelengths of

690nm, 744nm, 807nm and 832nm. Essenpreis et al [149] reported significantly lower variance of only ±9% at a single wavelength of 800nm, but on a much smaller group of only 6 healthy volunteers. The lower variance at the single 800nm wavelength may well have been due to the equal absorption coefficients of haemoglobin and oxy-haemoglobin at this wavelength, meaning that variance in tissue oxygen saturation did not affect the path length at this wavelength, and figures 6 and 7 (in the publication) would appear to support this hypothesis, although the small group in the Essenpreis study makes it difficult to draw definitive conclusions.

Errors in the order of 20% would appear to be sufficiently low to allow use of differential path length based NIRS in detection of dangerously low peripheral perfusion, or in long term trend based monitoring where absolute accuracy is not essential. However, Van der Zee et al [232] found an increase in both inter and intra subject variability with decreasing source to detector separation in all studied tissue types, with a particularly significant increase below 10mm source to detector separation, hypothesised to be partly due to the increasing relative inhomogeneity of the tissue layers as the distance was decreased (figure 2 in [232]). Intriguingly there is a considerably larger inter subject variability amongst the male volunteers in the Van Der Zee study, and although the total number of subjects is small it would appear that there is more variability amongst the male forearms, potentially indicating more variable muscle mass. It should be noted that the source and detector fibre bundles were repositioned for repeated studies of the path length on each volunteer, potentially leading to errors in the measurement of separation - Van Der Zee estimated this error as having a standard deviation of 2mm, and if this was indeed the case, then the majority of the increase in variability with decreased separation is easily explained by the increasing relative magnitude of the positioning error. Extrapolating the results from Van Der Zee to the larger sample group of Duncan et al, and assuming a 2mm standard deviation in the optode positioning, it would appear that the standard deviation at 5mm source to detector separation will increase only slightly to around ±20% or so.

From figure 2 in the publication by Van Der Zee et al it can be seen that there is a significant increase in the differential path length factor at shorter separations, particularly as the inter optode separation decreases to less than 20mm. Considering typical μ_{eff} in forearm tissue is of order $4cm^{-1}$ (and $\mu_{eff} = \sqrt{3\mu_a\mu_s'}$), it can be seen from equation 8.44 that as the separation d decreases to distances of order 1cm or less, the extra d dependent term in the expression starts to become significant, causing the DPF to increase. However, Van Der Zee et al hypothesise that the observed change is primarily due to the low numerical aperture of the source and detection fibre bundles. This

limited acceptance angle preferentially selects paths that pass deeper into the tissue at the source and detection regions, leading to an approximately constant path length increase. As the source to detector distance is decreased, this path length increase becomes more significant as a proportion of the overall path, leading to the increase in the path length factor. Van Der Zee also presents Monte-Carlo simulation results that suggest a decrease in effective path length factor as the source to detector distance is decreased, results that are apparently at odds with equation 8.44 derived by Boas. The fact that the model used in the paper by Van Der Zee et al used a layered tissue model may explain some of this apparent discrepancy, as equation 8.44 applies the diffusion model to a homogeneous semi-infinite slab.

Essenpreis et al [149] used a swept source pulsed titanium-sapphire laser to determine the wavelength dependence as well as the magnitude of the path length factor in detail. A wavelength scan consisted of 13 separate measurements made over a range of 740 to 840nm. From figure 4 in the publication, it can be seen that there is a consistently shaped wavelength dependence to the path length factor. The path length factor has a consistent downward trend with increasing wavelength, apparently due mainly to the decrease in μ_s' with wavelength. The data from Essenpreis are shown together with data from the BORL group at UCL published by Kohl ([148]) in figure 8.19. It can be seen that there is close agreement between the two data sets, and also a similar wavelength dependence to the pulsatile modulation depth based results produced by Kohl.

Due to the consistent path length factors reported by multiple authors, and the relatively low path length variability found within large groups of volunteers, it would appear that the differential path length model is suitable for use with the custom NIRS sensor hardware. It would seem reasonable to model a capillary refill as a time dependent change in the concentration of only two absorbers - haemoglobin and oxy-haemoglobin. If this simple model is used, it will be possible to process the raw five wavelength intensity data to path length corrected intensity changes and then apply a least squares fitting algorithm to obtain the changes in the concentration of these two absorbers with an overdetermined fit. Thus it will also be possible to obtain a measure of the fractional fit error in a similar manner to the methods used with the previous diffusion model optimisation. If the differential path model is applied to sample volunteer data, the magnitude of this fractional fit error will reveal how accurate the model is when used with the custom sensor and on the palm of the hand, rather than across the top of the forearm as in the UCL and Essenpreis studies.

It should be noted that such a fractional fit error is only suitable for determining if the shape of the candidate path length function is a good match to the data. It is not possible

to determine the magnitude of the path lengths by fitting this model to sample data, but if the wavelength dependence has a similar relationship to the published studies, then it would seem reasonable to conclude that the changes in source to detector spacing between the published UCL and Essenpreis data (which used 40mm separation), and the custom NIRS sensor (5mm separation) as well as the change in location on the subject (top of forearm to the Thenar eminence on the palm of the hand) do not render the published path length data invalid (or vice versa).

As it has been used in multiple studies and for the processing of data from commercial devices (the continuous wave Oxymon device [233] was developed at UCL before being commercialised by Artinis Medical B.V., Zetten, Netherlands), the UCL path length dataset as published by Kohl et al was chosen for use in the trial data processing algorithm. Unfortunately this dataset does not extend below 730nm wavelength, as can be seen in figure 8.19 (green stars). As can be seen from figure 7.12, the path length function needs to extend to at least as short as 610nm to cover the emission spectrum of all the LEDs used in the custom NIRS sensor. Unfortunately there are few published studies of typical path length factors at wavelengths of less than 700nm, so it is difficult to assess the accuracy of the estimated path lengths in this region, although the data from Duncan et al at 690nm suggests that the path length continues to increase below 700nm. The validity of the model in figure 8.19 to accurately predict the change in path length is cast into doubt by its inaccuracy around 950nm, and the fact that it would appear that other absorbers besides haemoglobins and water give rise to significant absorption in tissue. For these reasons it may seem appropriate to extend the UCL data from 730nm to 610nm using simple linear interpolation as a path estimate. In order to investigate the likely validity of such an approach, the diffusion modelling used in figure 8.19 was repeated and extended to a shorter wavelength range.

From figure 8.20 it can be seen that the increase in tissue absorption from both haemoglobin species below 700nm wavelength causes a decrease in the path length factor. Perhaps more significantly, there is a strong dependence on the tissue oxygen saturation, particularly around 650 to 700nm, a wavelength range that covers most of the emission from the 660nm LEDs on the NIRS sensor. This means that the path length will have a relatively high dependence on the tissue oxygenation, a situation that is likely to reduce the accuracy of the tissue molarities obtained via the path length method.

Due to the significant effects on path length from other absorbers in the tissue (other than haemoglobin and water), it is unclear how much of the decrease in path length seen in the model will be present in real tissue. However, due to the fact that the modelled path length decreases to well below that seen if the UCL dataset is extended by simple

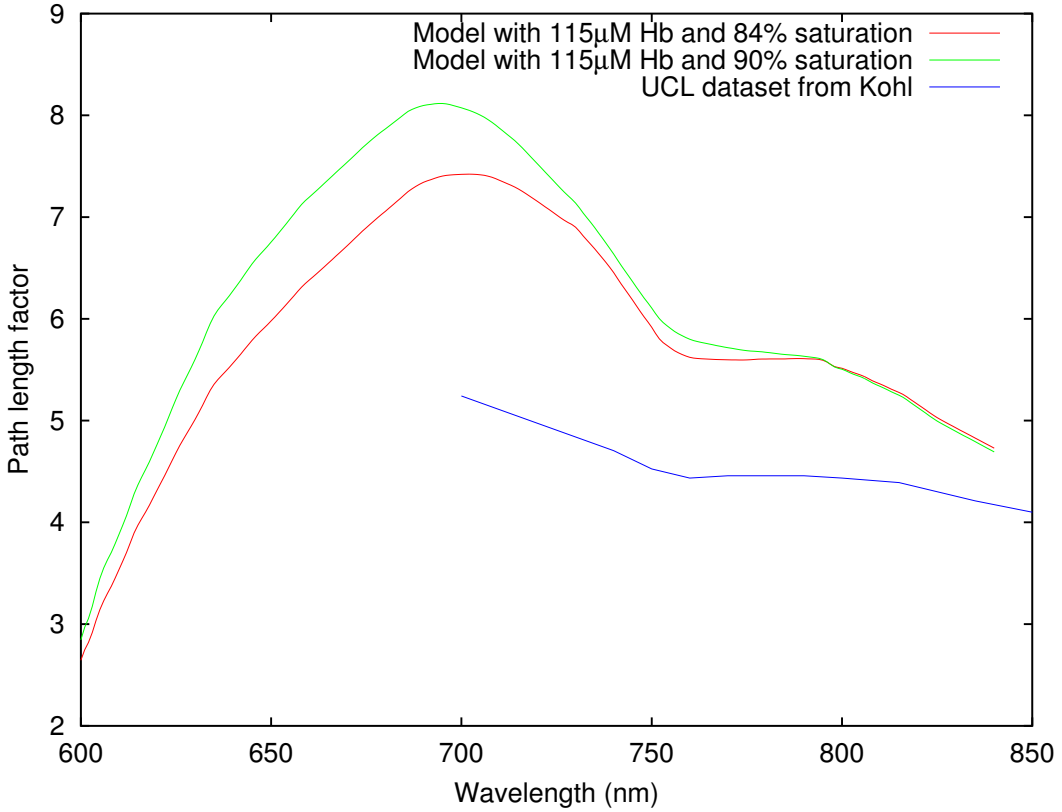


Figure 8.20: Path length factor computed from equation 8.43, repeat of previous plot at shorter wavelengths.

interpolation, it would seem reasonable to assume that the path length factor at 660nm wavelength will be in the order of 5 or less, rather than 6 or more as would be suggested by extrapolating.

One simple way to avoid this uncertainty entirely would be to use only the 4 larger wavelength datasets. With 4 wavelengths there are still sufficient degrees of freedom to solve for changes in Hb and HbO_2, and the path length factor at 660nm can be found using an iterative fitting process applied to the recorded datasets.

In order to calculate effective path lengths for each of the longer wavelength channels, a numerical integration process was used, similar to that used earlier to calculate effective absorptions for each LED type. Both the UCL path length and the LED emission spectra were first re-sampled to 1nm resolution using a cubic spline process, and the cumulative trapezoid numerical integration technique used to calculate the effective path length for each LED channel of the NIRS sensor.

It would be useful to be able to estimate the error in the Hb and HbO_2 molarity changes obtained using the differential path length model. As has already been discussed, obtaining an estimate of the likely path length uncertainly is possible from published studies,

but application to the custom made NIRS sensor is not trivial, as the source to detector separation is shorter than in most CW NIRS devices described in the literature. The variance in path length factor between individual subjects has already been discussed, but for a thorough error analysis of multiple channel data, the covariance matrix for the path length factors for the set of LED channels must be known. Equation 8.46 can then be applied to this covariance matrix. If the errors are assumed to be small, so a first order differential model can be applied first, followed by application of a change of basis operation:

$$\mathbf{Q}_{path} = \mathbf{A}^{\dagger}\mathbf{C}_{path}\mathbf{A}^{\dagger^{\top}}. \tag{8.53}$$

Where $\mathbf{C}_{path}$ is the covariance matrix for the vector path length factors, and $\mathbf{Q}_{fit}$ will be the covariance of the vector of absorber concentration changes. These two covariance matrices will be in relative rather than absolute units, meaning that equation 8.53 will give an absorber covariance matrix that tends to zero at the reference intensity. This is clearly not accurate, as amplifier noise, shot noise, and time variant optical coupling lead to absorber concentration errors. Thus for a full error model to be constructed, a separate fixed term covariance matrix $\mathbf{C}_{fixed}$ must also be computed:

$$\mathbf{Q}_{fixed} = \mathbf{A}^{\dagger}\mathbf{C}_{fixed}\mathbf{A}^{\dagger^{\top}}. \tag{8.54}$$

If the digaonal operator $Diag$ is introduced to project a vector onto a matrix diagonal, then these two error components can be combined to give the final covariance matrix estimate for the absorber concentration change vector in absolute units:

$$\mathbf{Q}_{fit} = Diag(\Delta\vec{\mu_a})\mathbf{Q}_{path}Diag(\Delta\vec{\mu_a}) + \mathbf{Q}_{fixed}. \tag{8.55}$$

Although no attempt at such a procedure was attempted in the Essenpreis paper [149], it would appear that there are sufficient data within this publication to calculate the covariance between deviation of the path length factor from the mean at different wavelengths. From figure 4(a) in [149] it can be seen that path length variability in the head is highly correlated across the wavelength range. This is to be expected, as variance in μ_s' and total blood volume fraction affects all wavelengths approximately equally. From figure 6, it can also be seen that in the adult forearm the variance in path length between subjects increases either side of 800nm. The data in figure 6 have been normalized relative to the 800nm path length, so the error bar at 800nm is expected to be zero, but if the errors were entirely random, then there would be no relationship between the variance and the difference in wavelength from 800nm. However a clear relationship

is seen, with the variance increasing with distance either side of 800nm. This can be seen to be caused by varying tissue saturation between the subjects; figure 8.20 clearly shows how the path length is extended below 800nm by increasing tissue saturation and decreased above this isobestic point.

Thus in conclusion it can be seen that there are three sources of the variance in the path length factor:

1. The variance due to variable tissue μ_s', tissue blood volume faction and concentration of other absorbers (e.g. Cytochrome oxidase). This effects the entire near infra-red wavelength range approximately equally.

2. Variance due to tissue oxygen saturation. This would appear to a less significant source of path length variability, at least from the relatively sparse set of available data.

3. Random path length fluctuations with no correlation between wavelengths. These are a very minor component of the path length variability in the published studies, but may be more significant in the custom NIRS sensor, where the LEDs are physically separate, and thus there is a different optical path at each wavelength.

As the likely nature of the path length variance in the custom NIRS sensor is unclear, it was decided to leave a full error analysis until an adequate dataset had been gathered. NIRS error analysis will be explored in more detail in Chapter 10.

In order to investigate the performance of the differential path algorithm further, a new dataset consisting of several minutes of repeated 5 second 3 PSI pressure applications followed by 10 second 0.2 PSI application intervals was recorded from the Thenar eminence region on the palm of a healthy adult volunteer. The raw intensity data were first normalized using a simple first order polynomial (i.e. mean) operation.

From figure 8.21 it can be seen that although the intensity during pressure application is relatively stable, there is a decrease in the size of the intensity decrease between applications as the recording progresses. The reason for this is unclear, there was no long term drift in the application pressure, as the linear actuator control system maintained a consistent 0.2 PSI between each pressure application period. One possible explanation is that some volumes of static blood (perhaps accumulated venous blood) were forced out of the pressure affected tissue region by the repeated pressure applications and did not return.

As extrapolation of the UCL path length dataset below 700nm wavelength is not possible with a high degree of certainty, only the data from the four longer wavelength channels

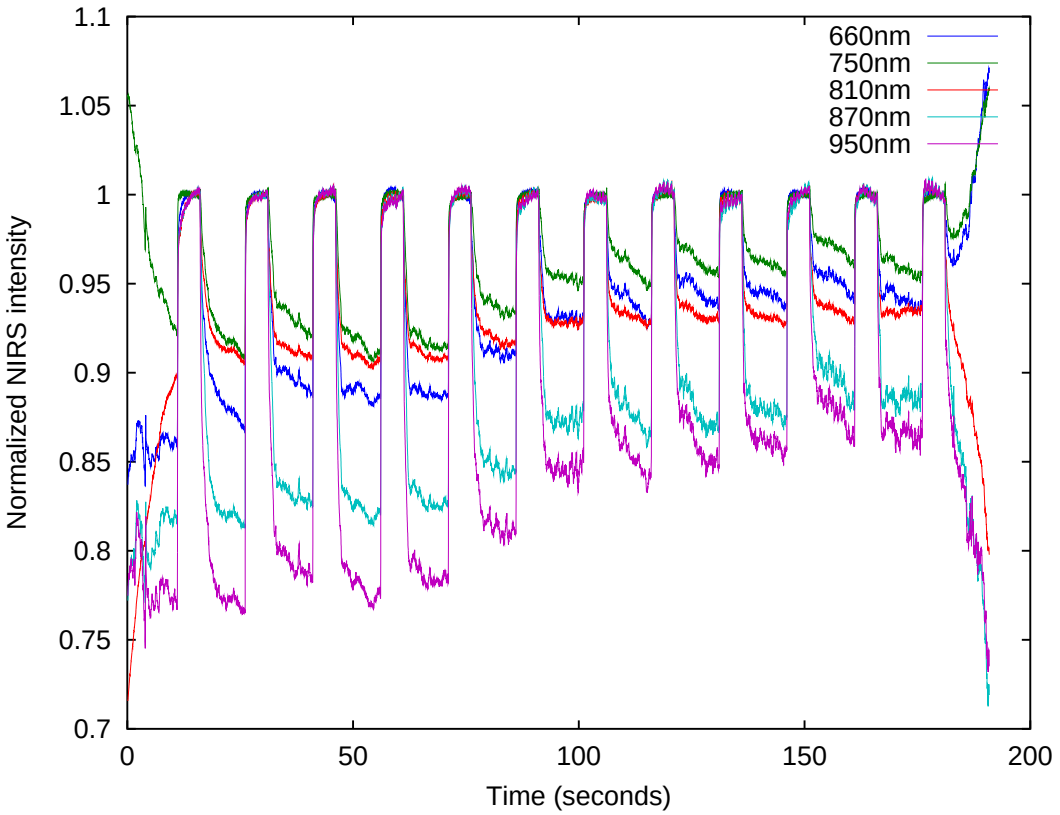


Figure 8.21: The normalized raw NIRS intensity data used for the fitting experiment.

was used in the fitting process (750nm, 810nm, 870nm, and 950nm).

Figure 8.22 shows the normalized data aligned to each 10 second period between pressure applications, together with the residual intensity shift remaining once the intensity change due to the fitted Hb and HbO_2 molarity changes is accounted for. It can be seen that the residual is relatively small at most of the wavelengths, but the 810nm residual is around 50%. Thus it would appear that the UCL path length data are not a good representation of the actual path lengths seen with the reflection mode NIRS sensor. This may be due either to a change in the path length function between the forearm and the palm of the hand, or to the design of the sensor itself, with shorter paths and widely spaced LEDs being used.

In the absence of more reliable path length data, it was decided to proceed with the differential path length approach, as it had the advantage of stability; unlike the diffusion model fitter attempted earlier there was no risk of the solution converging to an unphysical molarity.

Figure 8.23 plots the change in tissue Hb and HbO_2 molarities following each release of pressure, as derived by the differential path fitting algorithm. The molarity data have been aligned to each pressure release interval, and as it was found that there was an

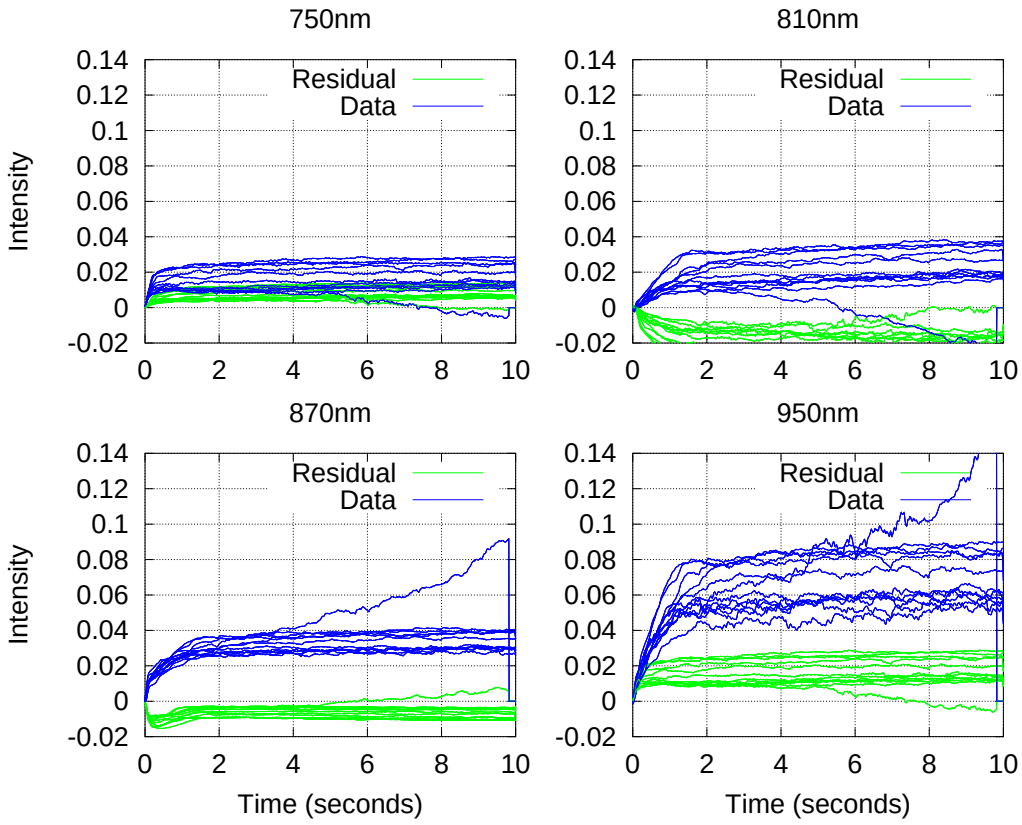


Figure 8.22: Oxy and deoxy haemoglobin fit to the raw NIRS intensity data in figure 8.21.

inconsistent molarity change immediately following each pressure release, the reference level (i.e. the zero molarity level) was set at a time of 300ms following the applied pressure falling below 0.5PSI.

It can be seen that the HbO_2 molarity increases with a approximate exponential curve following each pressure release, before falling very rapidly when the pressure is applied. The Hb molarity changes are considerably smaller, and Hb molarity change appears to roughly mirror HbO_2, decreasing after pressure is released. The Hb decrease after pressure release might be explained by Hb buildup within the tissue during the pressure application period when perfusion is halted, but this explanation would not lead to an immediate increase in Hb molarity seen as soon as the 3PSI pressure is applied.

Thus it would seem that an inaccurate path length function is the most likely explanation for the Hb molarity changes; if the blood volume change during and after 3PSI pressure application consists principally of oxygenated blood, i.e. a HbO_2 molarity change, then a small error in the shape of the path length function could lead to the Hb molarity change being inverted.

Looking at figure 8.23, it can be seen that although the total haemoglobin (sum of the Hb

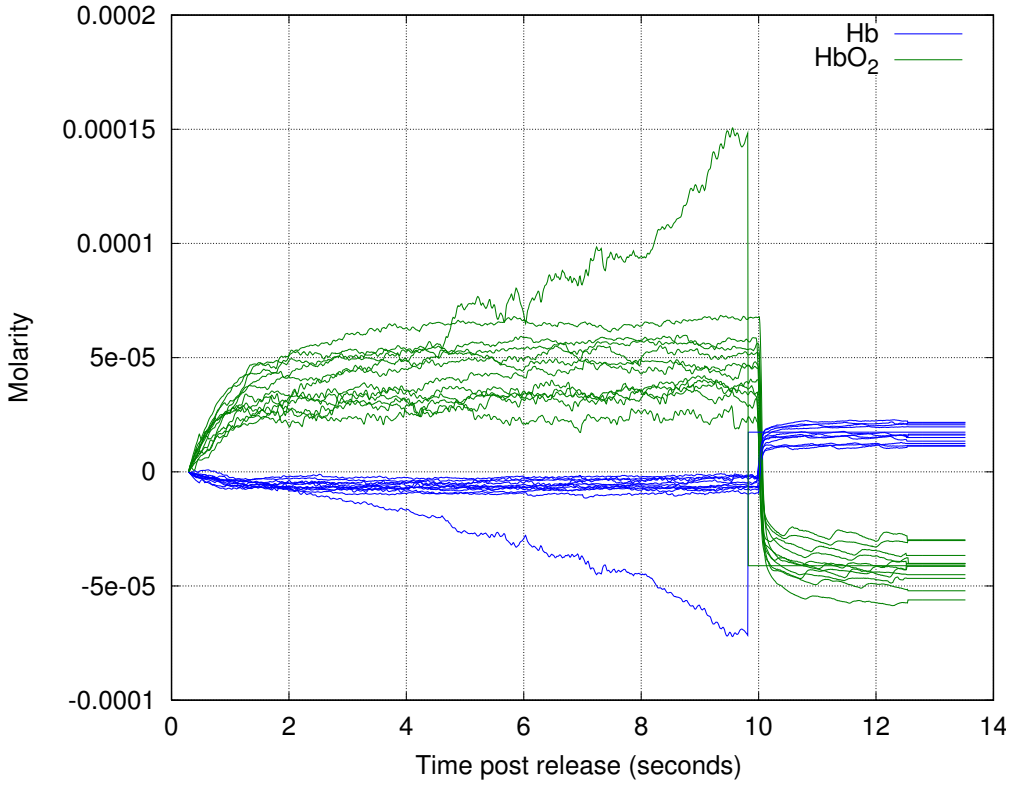


Figure 8.23: Oxy and deoxy haemoglobin fit to the raw NIRS intensity data in figure 8.21.

and HbO_2 molarities) increase changes between refills, the initial rate of molarity increase between approximately 300ms and 1300ms after the pressure release is somewhat more consistent at around $50\mu Ms^{-1}$. This corresponds to a volumetric perfusion of order 0.02ml/ml/s, a value more than an order of magnitude higher than typical values for the healthy adult palm found in the literature.

Thus it would seem reasonable to conclude from these initial results that the capillary refill does not provide a direct measurement of arterial perfusion. Blood inflowing through the capillaries and micro-vasculature from the surrounding tissue would appear a more likely source for the refilling blood flow, but this flow may still provide useful perfusion information, for example if the perfusion pressure (commonly defined as the pressure drop across the capillary bed) is known, then the perfusion will be a function of perfusion pressure and the capillary resistance, factors that will also influence the arterial perfusion.

For these reasons it was decided to further develop the automated capillary device. However, to allow investigation of the correlations between capillary refill and the volumetric perfusion from the arterial flow, it was decided to develop a venous occlusion NIRS technique using the custom sensor hardware in parallel.

Chapter 9

Development of a portable NIRS data logger with a pneumatic supply for controlled pressure application

In Chapter 7, a set of continuous wave NIRS techniques were identified as being amongst the most promising for practical non-invasive perfusion monitoring. In the previous chapter the development of a custom made LED based NIRS system was described. This system was then combined with an automated pressure application system to explore the potential of an automated capillary refill test in perfusion monitoring.

However, the resulting desk mounted device was a quite impractical solution, suitable for research but not clinical use. Moreover, it was only suitable for monitoring of the capillary refill, meaning that other techniques such as venous occlusion could not be explored without additional hardware. For these reasons a portable battery powered data logging device was developed in parallel with the desk mounted device. The NIRS system described in Chapter 7 was used as a basis for the new device, and it was decided that the addition of a controlled pneumatic supply, together with a data logging capability and battery power supply would permit stand alone operation.

9.1 Development of a portable continuous wave NIRS system

The "STM32F4-Discovery" development board used in the previous NIRS design is clearly not an ideal system for use in a portable device due to the relatively large size and lack of a battery based power supply. With this in mind it was decided to design a new custom PCB populated with a copy of the previous NIRS electronics along with a battery based power supply, charging circuit, on-board memory and interfaces for driving pneumatic components.

The STM32F4 micro-controller series is only available in packages of 64 pins or larger, and a small device was desired, so for this reason the micro-controller was changed and a 48pin STM32F103 used. This is a slower 72MHz ARM Cortex-M3 based device that lacks the FPU of the processor used in the bench device. However, it does have low current consumption both when running and in standby state with a real time clock enabled, ideal characteristics for a battery powered data logger, where the real time clock can be used to keep logged data organised. Benchmarking of the STM32F4 based system revealed that the processor utilisation was extremely low at less than 10% at all times, so it was also judged that the original processor was excessively powerful for the demodulation and pressure control tasks.

As design of a data logging system is somewhat complex, with a lot of potentially conflicting aims, it was decided to draw up a list of design goals that the data logging system should fulfil:

- A USB interface, ideally with driverless operation using the mass storage class for example.
- A simple means of battery recharging, such as via the USB interface (the USB standard supports a supply current of up to 500mA to an attached device).
- Enough onboard storage and battery capacity for multiple recording sessions, i.e. ten hours or more.
- A simple and user friendly way to set the onboard real time clock, i.e. one that does not require installation of driver software on the host PC.
- A wireless interface to allow data to be read in real time without cabling to the data logger. This also simplifies electrical safety.
- Compliance with relevant standards, for the Medical Device Directive 93/42 and its 2007 amendment. For hospital use as a research device CE is not required,

but documentation and EMC testing to an equivalent level simplifies the approval process for trials on patients within a hospital. In practice this means:

1. A single interface socket for all cables (i.e. USB lead and sensor cable), meaning a potentially unsafe configuration with the device connected to a PC via the USB interface with a patient sensor also attached is impossible.
2. All indicator LED colours conforming to their function according to relevant ISO standards (BS EN 60601 series, BS EN 60073).
3. Electrical isolation (in terms of leakage currents and Means Of Operator/Patient Protection - MOPP/MOOP) to relevant standards for a class 1 medical device (Section 3 of BS EN 60601-1:2006+A11:2011).
4. Photo-plethysmogram LED intensities assessed for eye and skin safety according to BS EN 62471.
5. A user manual describing correct and responsible operation of the device produced in accordance with Medical Devices Directive 93/42/EEC and Directive 2007/47/EC.

As it is widely used and supported, Bluetooth was selected for use as a radio link. Bluetooth PCB modules with CE certification are available off-the-shelf, meaning complex and expensive certification of the RF link is not required. A lithium polymer cell was used as a standalone power source, having a high energy density and a relatively high output voltage of between 3V and 4.2V from a single cell, simplifying the power supply circuit, and a μSD memory card was used as on-board storage for logged NIRS and pressure data. Unlike larger SD memory cards, all μSD cards support a simple SPI serial interface meaning only 4 pins of the STM32 micro-controller are required.

Figure 9.1 shows a symbolic system diagram of the design for the data logger electronics. Recharging lithium polymer cells is a relatively complex process of constant current followed by constant voltage trickle charging. Overcharging or using an excessive charge current can result in cell damage, presenting a fire hazard. For these reasons it was decided to use an off-the-shelf lithium polymer charging IC rather than the simpler approach of using the micro-controller to control charge current via a pulse width modulated transistor and current monitoring shunt resistor. An interface circuit was necessary to allow the micro-controller to be placed into or out of sleep mode and to determine the boot source when it awakes from sleep, as the micro-controller must not only enter and exit the shutdown state upon insertion/removal of the USB interface lead or the appropriate button sequence, but also identify the appropriate boot sequence of either USB device or data logging.

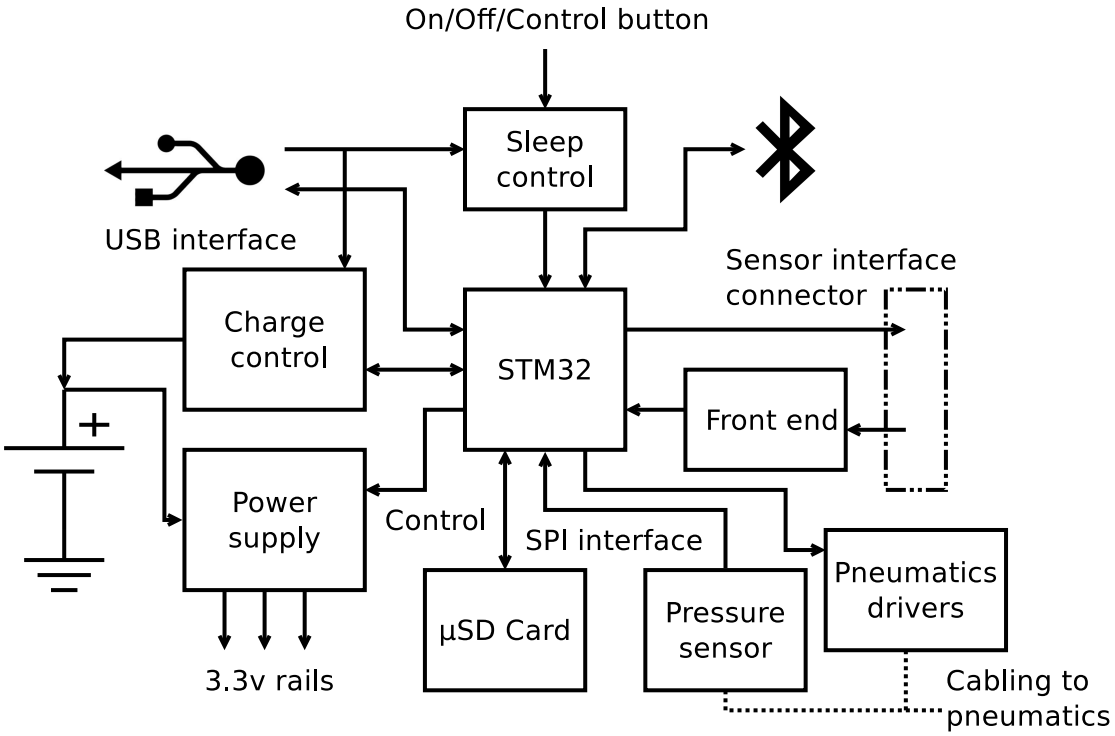


Figure 9.1: Proposed NIRS data logger system

9.1.1 Prototype data logger electronic design

A small quantity of prototype boards were designed and built, in order to test and debug the initial data logger design ideas. These boards were designed to be easy to modify, with the components spaced more widely than intended in the final production device, and with 0.1 inch headers used for connectors wherever possible. Nevertheless, the prototype boards were designed to meet applicable safety standards for hospital use, just as would be applied to any production device.

Use of a 3.3V supply rail was necessary in order to power a class 1 Bluetooth radio module (10mW output power giving operating ranges of tens of meters in an indoor environment) and use the single op-amp photodiode frontend design described in Chapter 7. As most sources suggest lithium polymer cells should be operated at an output voltage range from 4.2V (change termination voltage) to 3V (discharge termination voltage), it was decided to use a SEPIC (Single Ended Primary Inductance Converter) switching power supply in the first hardware prototype. As a SEPIC converter can step the input voltage up or down, the lithium polymer cell could then be used with optimum efficiency when its output voltage was higher than the 3.3V supply rail. The full recommended discharge cycle to 3V could also be completed as the converter transitioned to stepping up the cell voltage as it dropped below the 3.3V converter output.

Unfortunately all switching power supplies suffer from two potential problems; switching

noise is always present to some degree in the output, and they have relatively poor voltage regulation response to transient current loads. For these reasons it was decided to increase the output voltage of the SEPIC converter to 3.5V, and then use multiple low drop-out linear regulators to power the multiple 3.3V rails on the logger. By placing a large ($120\mu F$) tantalum capacitor between the 3.5V rail and ground it was possible for the linear regulators to adjust to changing current demand without the SEPIC converter having to respond immediately. This served not only to isolate the supply rails from switching noise, but also to isolate them from one another, so for example noise from the micro-controller digital core could not pass easily into the analogue front-end section via the supply rail. By using linear regulators with enable pins it was additionally possible to switch some of the 3.3V rails on and off using GPIO control to implement power saving schemes. Furthermore, the SEPIC - LDO topology also meant that a "pi" filter using ferrite suppression chokes and low ESR ceramic capacitors could be placed between the tantalum reservoir capacitor and the linear regulator for the analogue 3.3V supply, giving further noise reduction for the sensitive photodiode front end.

A third potential issue arises from the use of a switching power supply; the quiescent current is relatively high for switching supplies, meaning that the battery life will be low if the real time clock is to be permanently powered on the STM32 micro-controller. In order to avoid the occurrence of this situation, a low quiescent current 3.0V regulator was connected between the 3.3V supply rail to the STM32 micro-controller and the lithium polymer cell, allowing the STM32 to run from a 3.0V supply voltage derived directly from the lithium polymer cell when the SEPIC converter was disabled.

Figure 9.2 shows the schematic for the prototype power supply design. All the capacitors were ceramic (some of low ESR MLCC construction) except the $120\mu F$ tantalum reservoir capacitor. Additional decoupling capacitors (not shown) were placed close to the power pins of all ICs, of the appropriate sizes as recommended by the datasheets - $10\mu F$ capacitors were placed close to the μSD and Bluetooth module supply pins, and $4.7\mu F$ capacitors close to the STM32 VDD pins.

As the STM32 analogue supply rail must be present at all times, two transistors and a diode were used to supply either 2.3V to the AVDD pins (when the SMPS enable line was low), or alternatively connect AVDD directly to the analogue 3.3V supply rail.

A differential pressure sensor (26PC05SMT from Honeywell Inc.) was used for monitoring the pneumatic pressure. This outputs a millivolt level differential signal, so an Analogue Devices AD627R instrumentation amplifier was used to convert the output to a single ended voltage, and amplify it to a suitable level spanning the 0 to 3.3V range that could be read by ADC2 on the STM32 micro-controller.

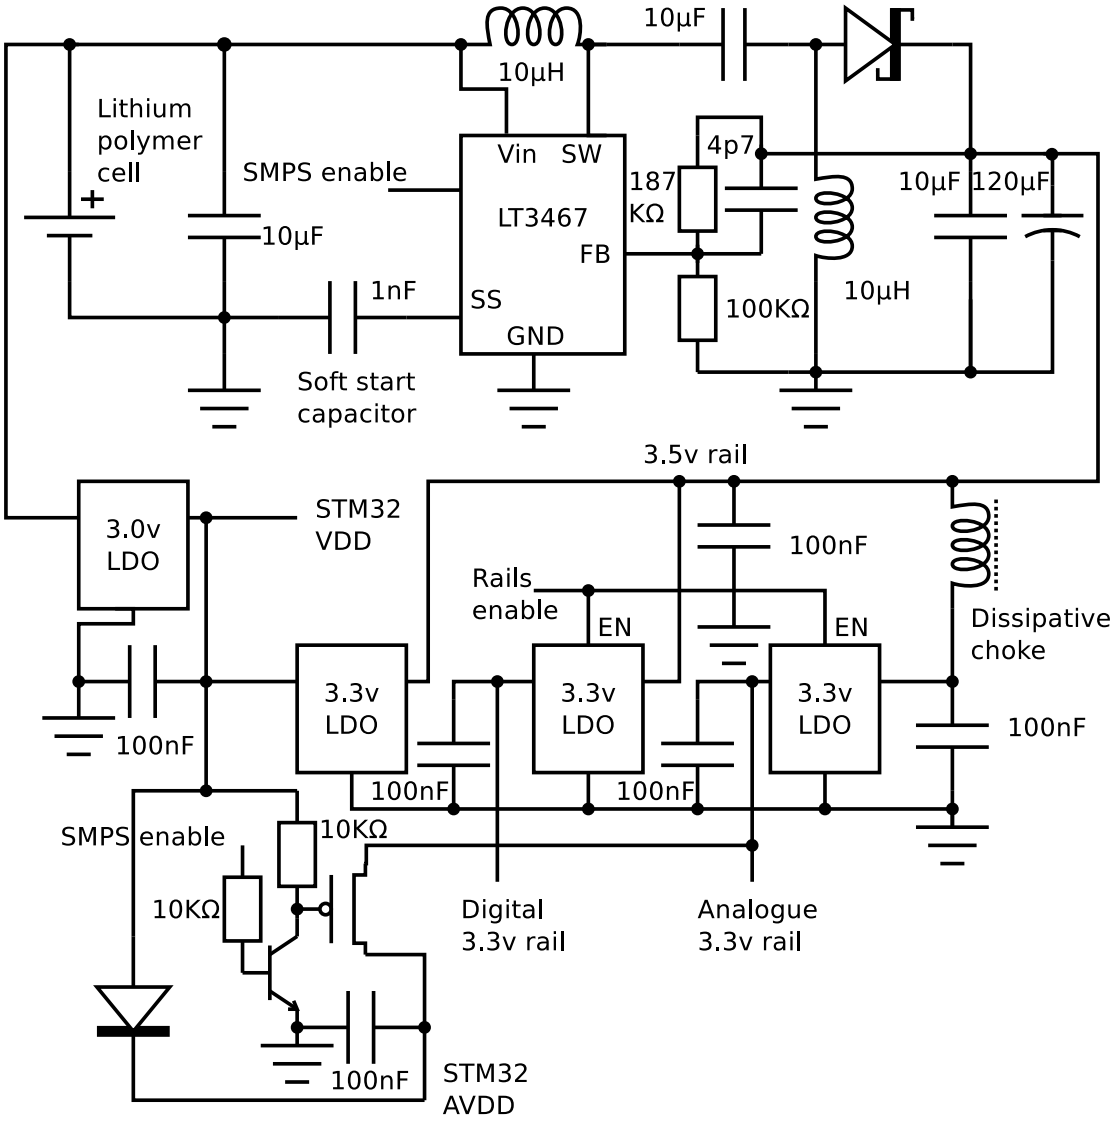


Figure 9.2: The SEPIC converter based data logger power supply.

The Maxim Inc. MAX1811 lithium polymer charger IC was used to charge the lithium polymer cell accurately from the USB 5V supply rail. This IC has a selectable 100mA or 500mA charging current, so to meet the USB current draw specification (100ma maximum current draw without an attached device) a pulldown resistor was used to select 100mA by default. Thus in the event of a firmware fault or a badly discharged battery with insufficient voltage to power the micro-controller (which had a GPIO pin connected to the charging current selection pin) the charger will default to 100mA charge current. By setting the pin high under firmware control once the USB bus is in the attached state, the 500mA charge current can be selected without violating the USB specification.

Two high current outputs for driving pneumatic components (i.e. air pumps and solenoid

valves) and accessible via a 0.1 inch pin header were also populated on the data logger PCB. However, the analogue front-end circuit described in Chapter 7 was not integrated into the main board, as it was decided that breaking out a 0.1 inch pin header to allow multiple analogue front end schematics to be benchmarked would allow the performance of the various designs to be investigated with sensors connected and used in realistic scenarios (e.g. under office lighting conditions).

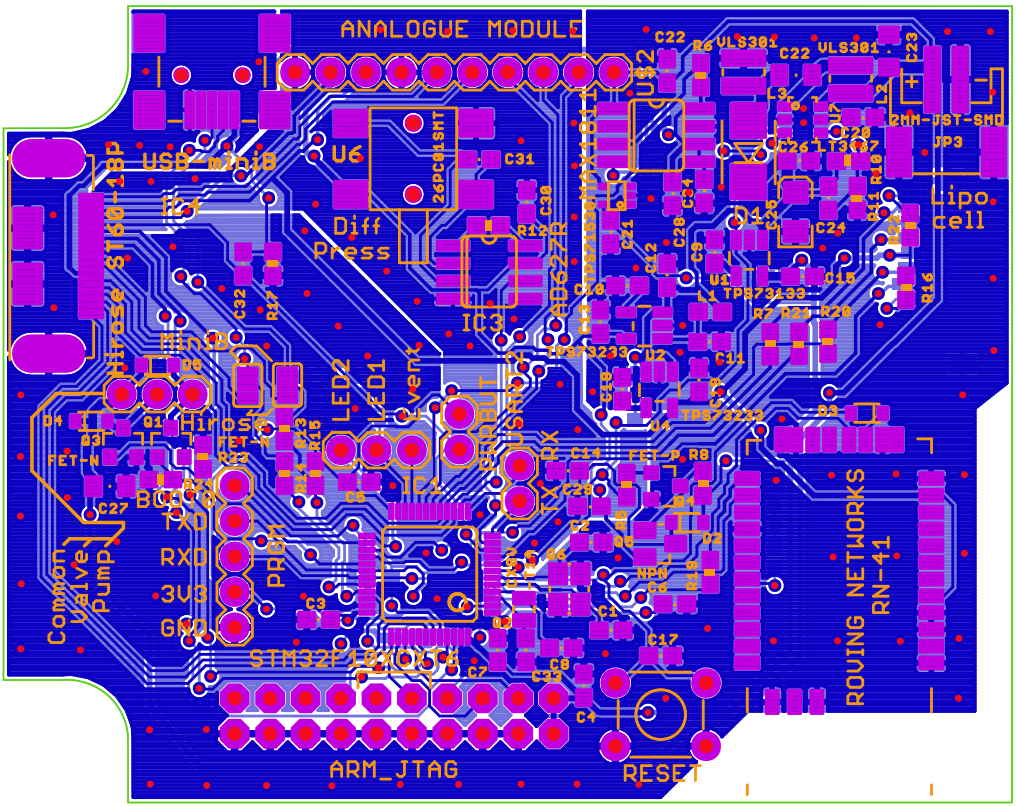


Figure 9.3: The PCB prototype designed for the NIRS data logger.

Figure 9.3 shows a rendering from the top side of the PCB design for the prototype data logger. It can be seen that the majority of the components are mounted on the top face of the board, with only the indicator LEDs and the μSD card placed on the bottom. The RN-41 class 1 Bluetooth module from MicroChip Inc. was used as the Bluetooth device. This has the advantages of being well documented and CE certified, simplifying the processes required to allow the data logger to be accepted into a hospital environment. A standard 20 pin, 0.1 inch ARM JTAG connector was placed below the STM32 micro-controller, allowing in system debugging, and an 18pin ST60 series connector from Hirose was used as the sensor interface connector.

9.1.2 Design of the data logger pneumatics

As outlined at the beginning of this chapter, the aim behind the controlled pressure pneumatic air supply was to allow both off-the-shelf blood pressure cuffs and pressure application devices for capillary refill studies to be driven by the data logger unit.

A large variety of miniature air pumps are available off-the-shelf.

- Piston pumps: these are mostly quite large devices designed for high pressure applications, with throttling being relatively difficult due to high stiction. For these reasons, piston pumps were judged unlikely to be suitable for use in the data logger.

- Rotary vane pumps: very small rotary vane pumps are widely available, and typically designed for use in gas sampling applications. As these pumps are reversible, it is possible to maintain a controlled pressure using only a single pump and an appropriate driver (e.g. a H bridge). However, miniature rotary vane pumps capable of generating pressures in excess of around 1PSI are less common, and higher pressure pumps typically use more expensive vane materials such as graphite vanes with a metal body, rather than polymers such as PTFE and Delrin (PTFE vanes and a Delrin casing being common). A number of gas sampling pumps were purchased to evaluate their feasibility, and it was found that flow rates dropped precipitously once the differential pressure was above 1PSI, making them unsuitable for use in either blood pressure cuff inflation or capillary refill sensors.

- Diaphragm pumps: these pumps are widely available, and are used in a wide variety of commercial products. Small (maximum dimension $\sim$5cm, mass $<$100gram) diaphragm pumps capable of pumping to pressures in the range of 10PSI or more with flow rates in the range of 1L/minute are commonly found in devices such as blood pressure meters and aquarium oxygenators. These pumps are not reversible, so if a diaphragm pump is to be used in a controlled pressure application system it will be necessary to use a solenoid valve to allow air to be vented from the system.

A diaphragm pump was selected for use with the data logger system due to the availability of suitable pumps from multiple manufacturers, and the high pressures and flow rates that can be achieved at a low cost. Figure 9.4 shows a schematic view of the pneumatic system. The chosen design is identical to that used in many commercial blood pressure meters, for example a tear-down of the Omron M2 blood pressure meter revealed an identical pneumatic system, and tear-down information published online reveals other off-the-shelf blood pressure monitors also using this design.

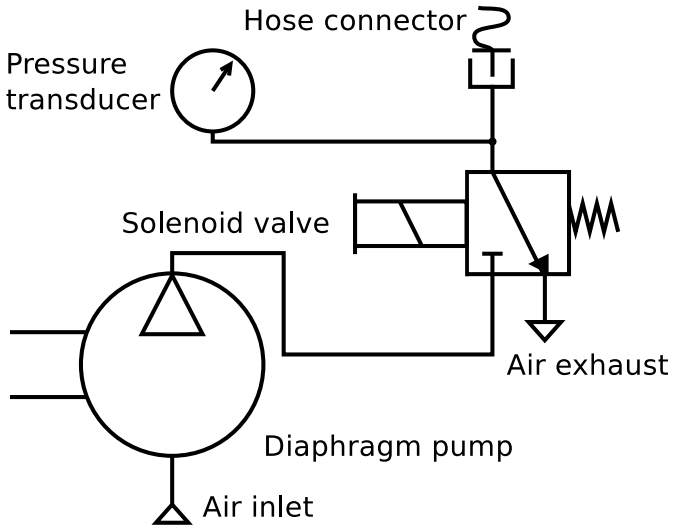


Figure 9.4: The pneumatic system designed for the NIRS data logger.

The diaphragm pump and solenoid valve were sourced from Parker-Hanniffin corporation (the T5-1IC-03-1EEB pump and 912-000007-005 2 way solenoid valve), and a blood pressure cuff connector from an Omron M2 blood pressure monitor was used as the air hose connector. It was found that the Omron pressure cuff connector was compatible with a wide selection of off-the-shelf blood pressure cuffs. Thus use of the Omron connector resulted in the additional advantage that off-the-shelf arm cuffs could be connected to the data logger unit for use in venous occlusion for example.

Figure 9.5 shows a top view of the partially assembled data logger unit. The Omron air hose connector can be seen below the main PCB (grey plastic). A 120 by 60 by 30mm high polycarbonate enclosure was used, with slots and holes milled for the connectors and the on/off/control button.

Several prototype front end modules were assembled for use with the prototype data logger electronics. The main PCB was designed to be capable of interfacing with dual photo-diode NIRS sensors, potentially allowing sensors similar to those described by Heuber et al [141] to be used with the base unit. Both of the front end designs explored in Chapter 7 were prototyped using separate plug in module designs. It was found that there was no detectable difference in signal to noise ratio between the two designs, with the noise being considerably higher than both the predicted value and those seen with the F4-Discovery based NIRS system. Benchmarking of the STM32F1 ADC confirmed that the effective number of ADC bits was in the region of 11.5, matching the datasheet values, so the only difference remaining between the data logger system and the F4-Discovery was the method of LED drive; the F4-Discovery used a FET connected to the appropriate PWM output from the micro-controller to drive each LED channel, but in the prototype data logger electronics the LED channels were driven directly from the

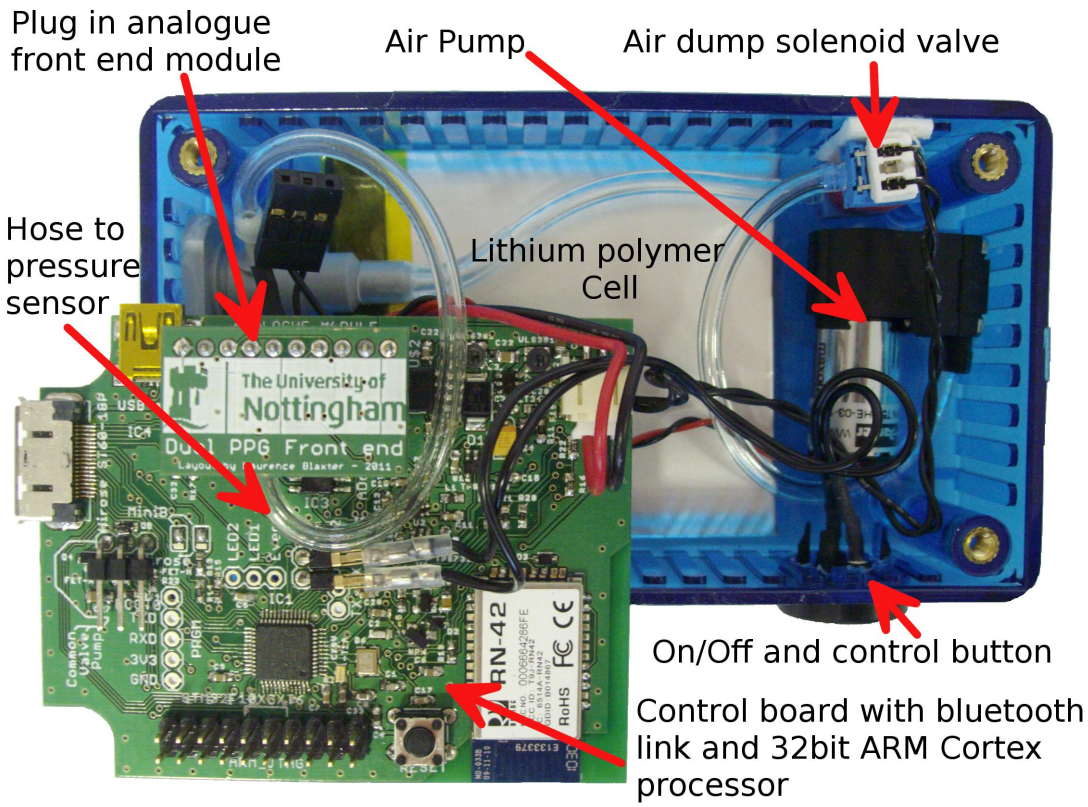


Figure 9.5: The NIRS data logger, partially assembled view.

micro-controller PWM outputs.

It was later found that driving the LED output channels with a high (20ma) current directly from the GPIO lead to significant fluctuations in LED current depending on the logic state of the other pins on the GPIO bank (see for figure 9.10 for an example of this problem). For this reason it was decided to use P-channel MOSFET transistors to drive the LEDs on the production version of the NIRS data logger system.

Figure 9.6 shows an external view of the assembled data logger system with the hose assembly of a prototype capillary refill sensor plugged in. The PCB was mounted so that the "bottom" side was facing upwards inside the enclosure, meaning that the indicator LEDs were visible through the enclosure. The μSD card can also be seen at the left hand side of the PCB face. The USB cable connection in the side of the unit was added purely for ease of firmware development, and means that the device does not meet electrical isolation standards for patient trials, as it would be possible for the USB cable to be connected whilst a patient applied sensor was also connected to the device.

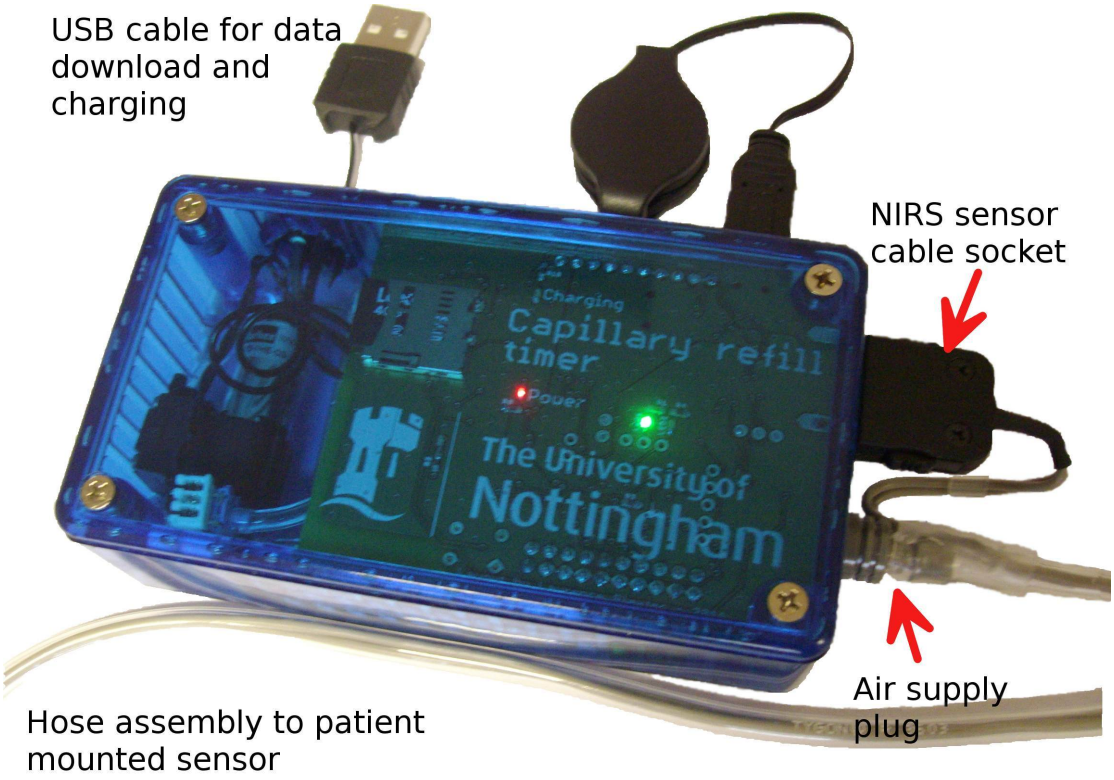


Figure 9.6: The assembled NIRS data logger with capillary refill sensor hose prototype connected.

9.1.3 Data logger firmware development

The data logger firmware drew heavily on the F4-Discovery based NIRS system, but a number of changes were made to the spectroscopy DSP to tailor it to the STM32F1 series micro-controller. The ChibiOS RTOS was not used for the data logger firmware, and interrupt driven code was instead written to use the peripheral library from ST microelectronics and the CMSIS abstraction layer from ARM. This gave more flexibility to use power saving modes and DMA than the ChibiOS hardware abstraction layer would allow.

The output sample rate from the NIRS DSP demodulator was reduced from 116Hz to approximately 62Hz. Combined with a lower ADC sample rate and a reduction in the number of orthogonal modulation channels from 5 to 3, this resulted in a significant reduction in the number of DSP operations per second, offsetting the lower clock speed of the STM32F1 micro-controller and the reduced DSP performance of the Cortex-M3 core. An orthogonal PWM scheme similar to that in figure 7.11 was used to drive the LED channels, with the frequency offset between channels equal to the new reduced DSP output sample rate.

An interleaved DMA transfer technique was developed in order to reach the highest possible transfer speed when transferring files over the USB interface. The USB interface on the STM32F1 has a dedicated SRAM memory area, but the DMA transfers were made via a separate area of main SRAM memory to allow both the USB interface and the SPI bus to the μSD card to be operated simultaneously without data corruption due to simultaneous access. The USB peripheral was configured with three endpoints, a control, data in and a data output; as described in the mass storage class specification. This resulted in an average transfer rate of approximately 750kBps when reading or writing large files. The theoretical throughput of the 12Mbps USB interface is in the region of 1MBps once protocol overhead and framing is accounted for, so there is clearly some potential for this interleaved DMA technique to be further improved.

In order to allow the real time clock hardware on the micro-controller to be set in a user friendly manner without the need to install driver software, a technique was developed to use the time stamp meta data of a text file (with filename "time.txt") saved onto the data logger from a host computer to set the micro-controller clock automatically. Thus, provided that the host computer had its system time set correctly, the only action required of a user to reset the clock was to save a blank file to the device.

The open source FatFS library was used to provide a simple API for all file and directory read and write operations on the logger, with the μSD card formatted using the FAT32 file system, which is supported by all main operating systems that might be encountered in use on a host computer. The FatFS library was only required for data logging operations, and not for USB interfacing, as the USB mass storage class requires only low level access to storage devices (using legacy SCSI commands sent to the control endpoint on the attached device).

The ARM Cortex system timer was used to trigger an interrupt service routine at 100Hz, from which time critical tasks such as the pneumatic pressure control and interpretation of button press sequences for "event" labelling of times of interest in the datafile were triggered. The ADC DMA half completion interrupt was used to trigger the NIRS DSP routine, which ran inside the interrupt service routine, rather than in its own thread as was the case on the F4-Discovery based system.

The gate of the diaphragm pump drive MOSFET was driven using a PWM output of the STM32F1. In order to reduce the impact of interference, the drive frequency was set as the third harmonic of the central NIRS modulation frequency, making the PWM pump frequency equal to approximately 36kHz. By making the modulation frequency equal to a harmonic, any crosstalk from the pump or the pump drive electronics to the analogue front end appeared as a constant DC offset on the demodulated intensity channels, rather

than as a rippling intensity term as would be the case at any frequency that was not an integer ratio of the NIRS modulation centre frequency. PID (Proportional Integral Derivative) control of the PWM duty cycle was used to control the diaphragm pump rotation rate in order to maintain a setpoint pressure, and for capillary refill studies the pump PWM was disabled and the solenoid valve opened during the zero applied pressure periods. The integral term of the PID control loop was also disabled during these zero pressure periods to prevent problems of so called "integral windup", where the I term increases out of control.

The Honeywell differential pressure sensor was calibrated using a two point linear model, as the datasheet specification gave a typical linearity of 0.5% and repeatability (i.e. hysteresis) of 0.2%. A manometer (Extech HD700, FLIR Systems Inc, NH USA) was used together with a syringe (for air supply) and water manometer (air accumulator) for this calibration procedure.

In practice it was found that the measured pressure suffered from very strong oscillation due to the fluctuating airflow into the hose assembly from the diaphragm pump, which only pumped air into the hose for less than one half of each rotation period, and suffered from significant leakage back through it's built in flap valves at other points in its rotation. The proportional and integral terms of the PID controller were not significantly adversely affected by this, but the very significant rates of change of the measured pressure resulted in the PID controller output being dominated by a derivative term that oscillated with the rotation of the diaphragm pump. This caused the pressure controller to become highly unstable.

A low pass filter was implemented to remove the pump noise, but it was found that filter lag was liable to cause the PID controller to become unstable. Figure 9.7 compares the results obtained with no filtering, a low pass Butterworth FIR filter, and a root-raised-cosine FIR filter that was further modified to reduce lag.

From the figure it can be seen that the control loop behaves chaotically in the absence of any pressure filtering. With the Butterworth filter applied there is excessive lag, causing the feedback to become unstable and the pressure controller oscillates heavily with poor pressure control. It was found that a root raised cosine filter with a relatively high transition frequency was sufficient to remove enough of the pressure noise for stable PID loop operation. In order to further reduce lag from the filter, the root-raised-cosine FIR table was asymmetrically cropped, reducing the time lag between incoming pressure data and the peak of the filter response to 70 milliseconds, and giving the comparatively stable pressure seen in the dotted green plot line. It was found that the performance of the low lag filter and PID control improved as the volume of dead space in the attached

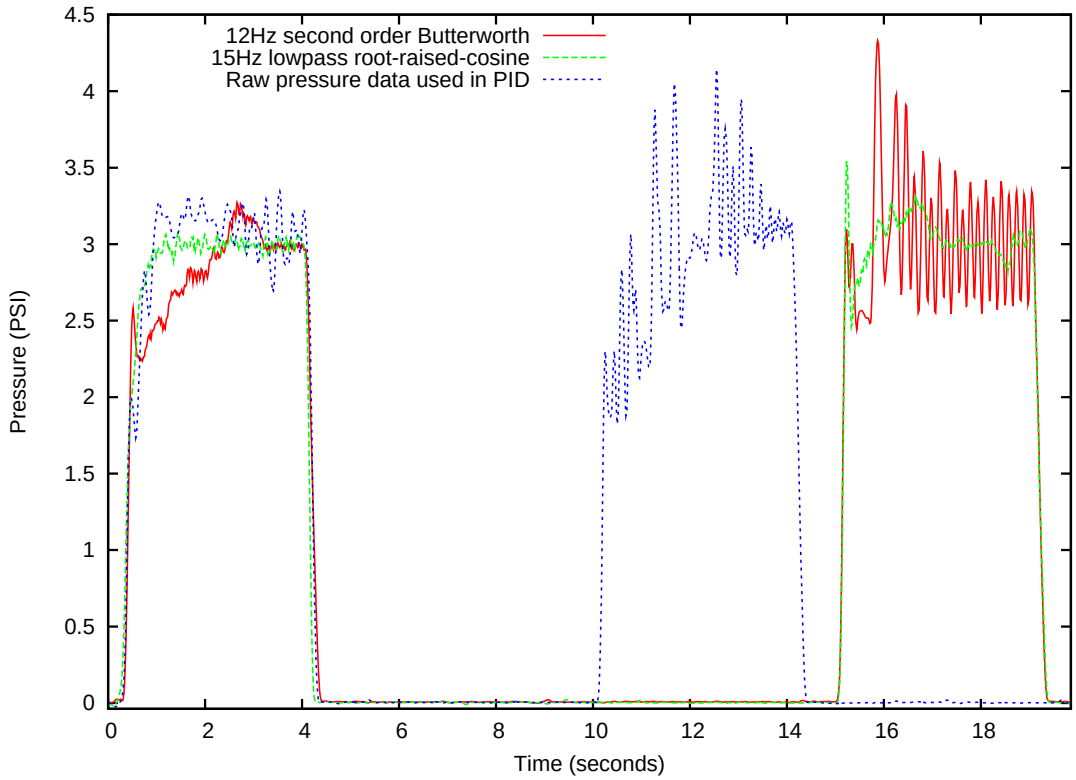


Figure 9.7: Pressure control loop stability for two filter schemes and with no filtering.

air hose increased. The data in figure 9.7 were recorded with only a 20cm manometer hose connected, but it was later found that pressure stability increased with 1 meter long sensor hoses connected to the device, at the expense of a small increase in pressure rise and fall time. This would seem likely to be due to the slight increase in response time of the system meaning the control loop had more time in which to respond in order to avoid overshooting.

9.2 Development of a pneumatic capillary refill sensor

By using either an off-the-shelf pulse oximetry sensor, or a modified copy of the sensor seen in figure 7.15, the portable data logger system could easily be used for venous occlusion and photo-plethysmography studies. However, in order to replicate the earlier capillary refill studies with hardware that was suitable for use within a hospital environment (i.e. that met appropriate health and safety and EMI/EMC standards), it was decided to develop a pneumatically actuated version of the NIRS automated capillary refill sensor.

9.2.1 Prototype pneumatic sensor

As an initial proof of concept, a prototype sensor was assembled using a trans-cutaneous oxygen probe attachment ring. A 1m length of bio-compatible (to ISO10993) "Tygon" PCV hose (sourced from Saint-Gobain plastics) was used together with commercial microphone cable passed down it's centre to drive the LEDs and interface with the photodiode. A connector from an Omron blood pressure cuff was reused to allow the data logger unit pneumatics to control the air pressure in the centre of the attachment ring, applying force directly to the skin surface of the volunteer. However, the safety of this approach is in doubt; for example the OSHA regulatory body (US Occupational Health and Safety Authority) states in its guidelines for workplace compressed air safety that air with a pressure greater than 10PSI carries a risk of air embolism if in direct skin surface contact. The Parker diaphragm pump used in the base unit has a maximum pumping pressure specification of 7PSI, which was considered too close to this 10PSI limit. For this reason it was determined that direct application of air to the skin surface would be unsuitable for a production sensor, and that a flexible membrane should be used between the air and the skin surface.

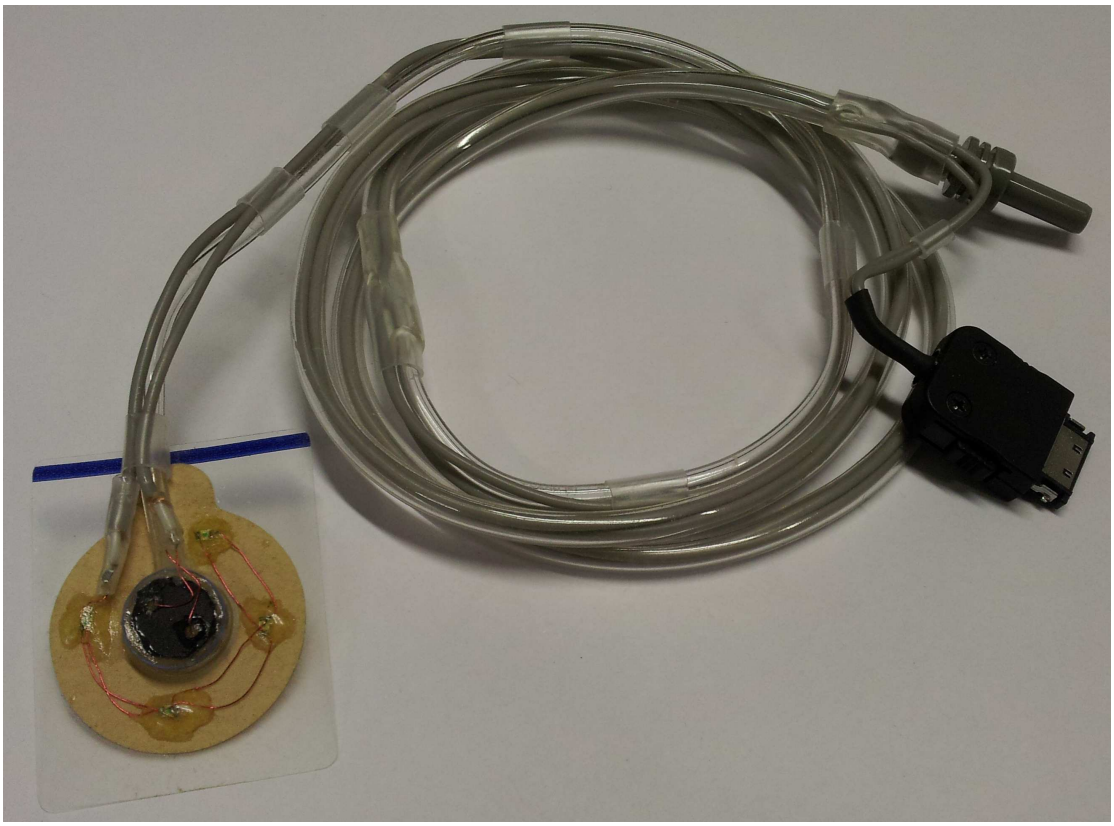


Figure 9.8: The proof of concept pneumatic capillary refill sensor prototype.

Figure 9.8 shows the sensor prototype together with the hose/cable assembly for attachment to the data logger base unit.

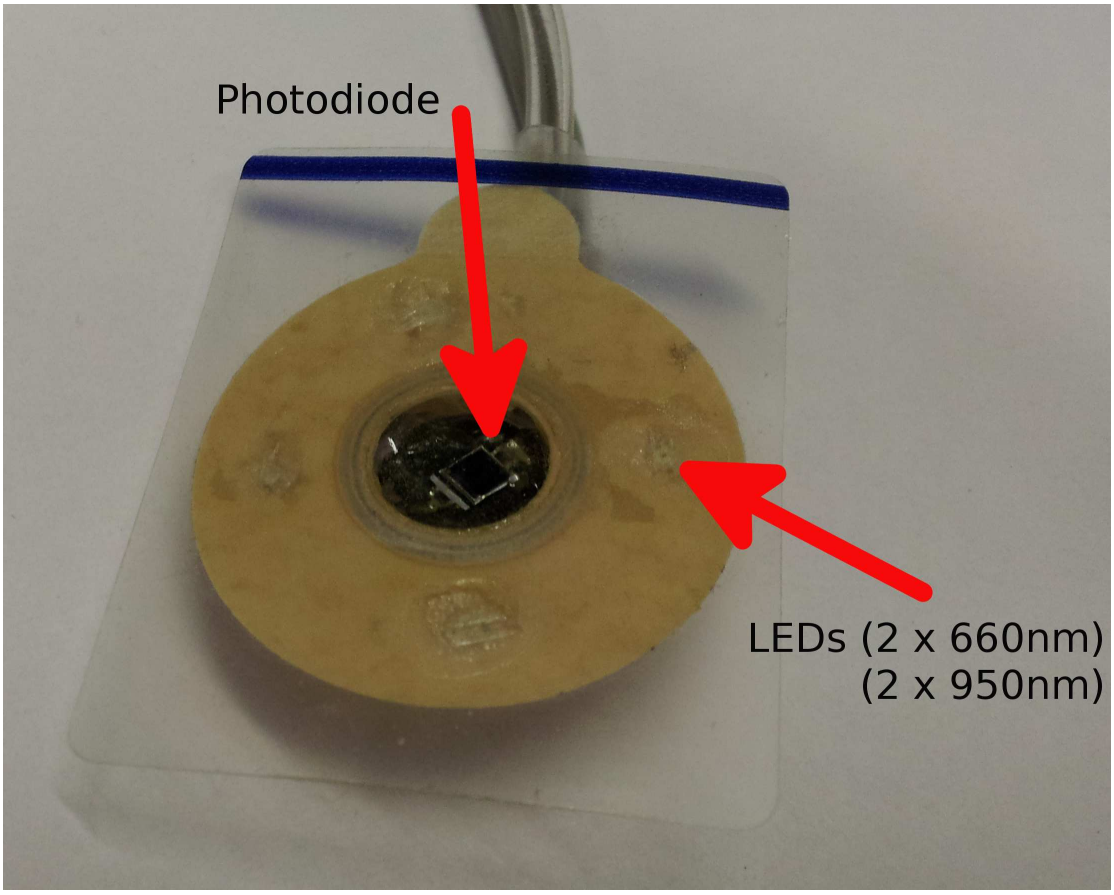


Figure 9.9: The bottom face of the pneumatic capillary refill sensor prototype.

Figure 9.9 shows a labelled view of the bottom face of the sensor. Source to detector separation is approximately 6mm, although the shortest possible optical path is approximately 4mm when exiting the skin surface at the inner edge of the attachment ring. In order to reuse the sensor the transparent plastic release liner was left in place and adhesive attachment rings designed for use with laser Doppler flow meters were used to attach the device to the skin surface of healthy adult volunteers.

Figure 9.10 shows normalized intensity data recorded from the dual wavelength prototype sensor. It can be seen that there is an approximately linear decrease in intensity over a 4 second period following the start of the pressure release. There is also a larger change in the red (660nm) channel than in the near infra-red (950nm) channel, suggesting that the refilling blood is de-oxygenated. However, considering the low accuracy of the LED placement and the possibility of light piping via internal refraction within the transparent plastic release liner and attachment ring, it is difficult to draw accurate conclusions from the data, other than that blanching combined with NIRS is feasible with such a device.

The two near infra-red LEDs used in the sensor prototype were driven from the same GPIO bank as the solenoid valve on the STM32F1 micro-controller. It was found that

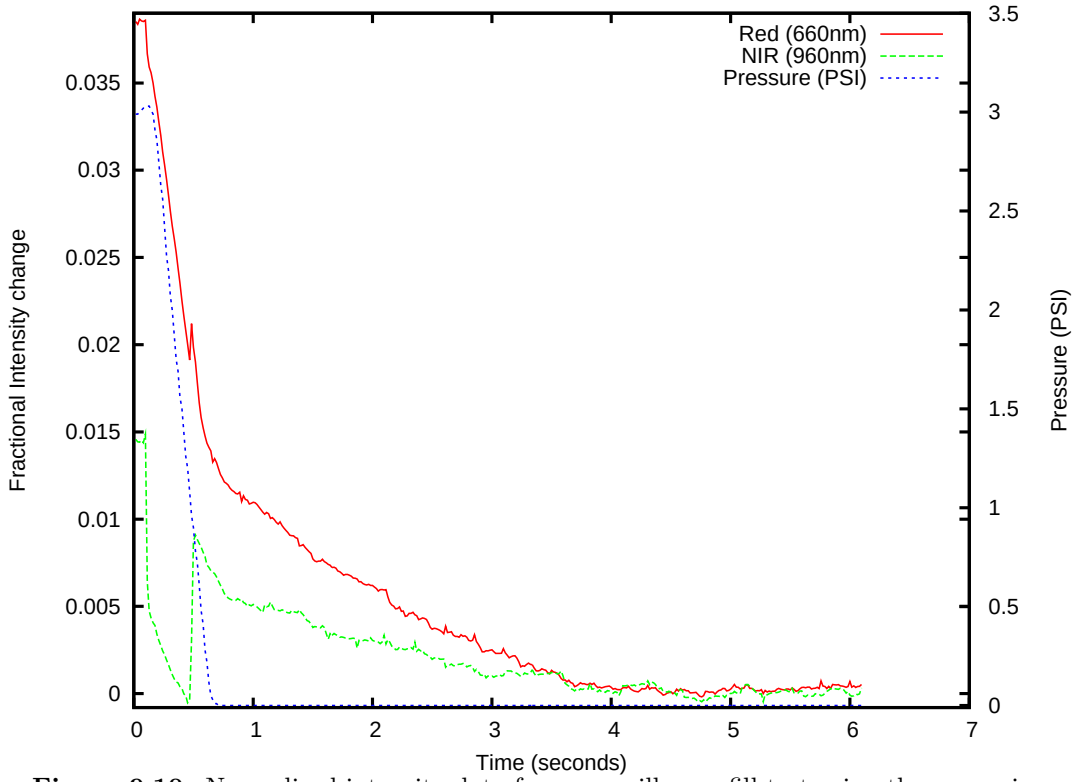


Figure 9.10: Normalized intensity data from a capillary refill test using the sensor in figure 9.9.

this resulted in a significant decrease in LED intensity when the solenoid was energized, even though no significant current passed into the gate of the solenoid driver FET. This effect can be seen occurring on the near infra-red intensity channel during the first 400 milliseconds of the pressure release, this being the time for which the solenoid valve coil was energized.

During further volunteer tests it was found that there was significant light leakage directly from around the LEDs (through the transparent epoxy resin used to encapsulate them) and into the transparent plastic of the attachment ring around the photodiode. For this reason a second prototype sensor was assembled using a black light absorbing epoxy for encapsulation of the LEDs. A temperature sensor was also added to the second prototype to allow an approximate measure of skin surface temperature to be made.

As the I2C bus interface of the STM32F1 micro-controller was broken out through the 18 pin connector on the data logger PCB, it was decided to use an I2C temperature sensor to avoid the calibration problems inherent in an analogue design. The TMP102 sensor from Texas Instruments was used for this purpose as it is available in a small SC-70 surface mount package.

From the example recording in figure 9.11, it can be seen that the sensor was slowly heated from the 25°C room temperature after being adhered to the palm of the vol-

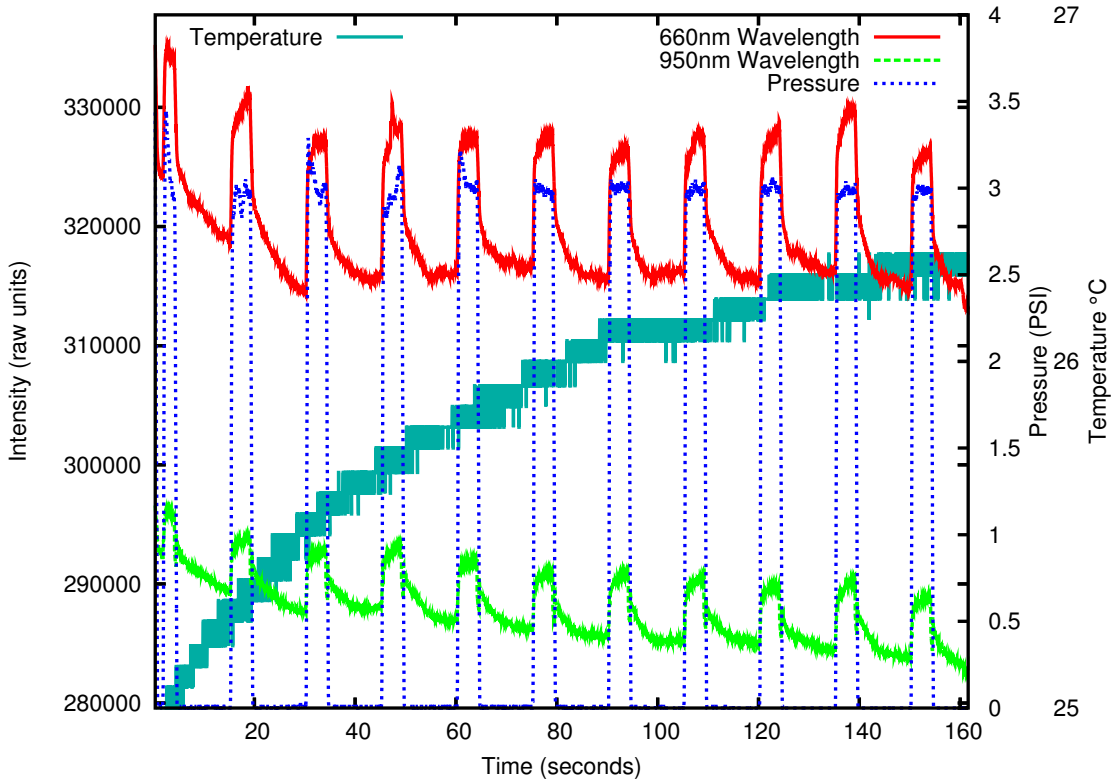


Figure 9.11: Raw intensity data from a capillary refill test using the second sensor, applied pressure and temperature are also plotted.

unteer's hand at 0 seconds. The response is relatively slow, with a time constant of order 1 minute or more, and it is unclear to what extent heat loss from the sensor to the surrounding air biases the measured temperature. The skin surface temperature would be expected to be significantly below core body temperature, making it difficult to draw conclusions here, but a temperature sensor placed closer to the skin surface (the TMP102 was adhered to the upper face of the sensor using epoxy resin) and with an insulating material placed over it to give isolation from the surrounding room temperature air would appear to be a more suitable solution to this problem.

It can also be seen from the figure that there is of order 1 least significant bit of noise from the temperature sensor. This is useful as it means that oversampling methods may be used to increase the effective resolution of the temperature sensor beyond the 0.0625°C resolution of the raw data. The apparent noise in the raw optical intensities plotted in figure 9.11 is mainly due to a strong pulsatile signal, although this is not clear in the plot due to the compressed time axis.

A final point of note in figure 9.11 is that the air pressure is clearly more stable than in figure 9.7. This is believed to be a result of the larger air volume within the sensor hose and sensor than within the short hose to the manometer used in the figure 9.7 recordings.

9.2.2 Development of a production ready pneumatic sensor

The prototype sensor based around a trans-cutaneous oxygen sensor attachment ring is clearly an impractical sensor design, since it does not meet bio-compatibility, electrical safety or compressed air safety standards, and is easily damaged.

For these reasons it was decided to develop a more suitable pneumatic capillary refill sensor. Based on the experience with the prototype sensor, a list of design goals was drawn up for the new proposed sensor.

- A bio-compatible flexible membrane material must be present between the pressurized air and patient/volunteer's skin.
- All external sensor materials that could contact the patient's/volunteer's skin must be bio-compatible to the ISO10993 standard.
- Electrical safety must fall under the applicable parts of the EN60601-1 (IEC60601) standard. It was decided that the patient applied part of the device (i.e. the sensor) should be designed to apply protection to a potential of 1kV, or 2 MOPP (Means of Patient Protection) as defined by the EN60601-1 standard (as the device is not mains powered. In the case of a mains powered device the voltage increases significantly to 3kV). This simplifies the approval process for hospital use, but means relatively strict electrical isolation design must be adhered to.
- The sensor must be sufficiently robust for sterilization/cleaning and reuse to be feasible. Although a disposable sensor would be a practical solution for commercial production, for medical trial use a reusable device would be much more practical.
- The sensor materials must be light absorbent to prevent light piping through the sensor directly from the LEDs to the photodiode. This was a significant problem in the prototype sensor.
- The optical path (source to detector separation) should be equal between all the LEDs and the photodiode. From the experience with the linear actuator based refill sensor it was decided to place the LEDs as closely as possible to one another in order that the optical paths between each LED and the photodiode should pass through as similar a region of tissue as possible.
- The light loss due to the refractive index mismatch between the sensor and the skin surface should not vary with the applied pressure. This problem was solved using a contact gel with the linear actuator device, but contact gel is not a very practical solution for a production device, and other solutions such as a skin adhesive coating

beneath the optodes may be more appropriate. Alternatively an air gap could be maintained between the LEDs and photodiode and the skin surface, resulting in a constant intensity loss at the skin surface interface, meaning there would be no opportunity for motion artefacts in the intensity data as a result of the changing applied pressure.

Several potential sensor designs were drawn up based on this list of requirements. The first sensor design that was investigated used a flexible, transparent, silicone rubber membrane to apply pressure directly to the skin surface. A cluster of LEDs and a photodiode were adhered to this membrane using Room Temperature Vulcanising (RTV) silicone rubber, and a polyamide flexible PCB was used to mount the sensor electronics and cabling to the data logger connector.

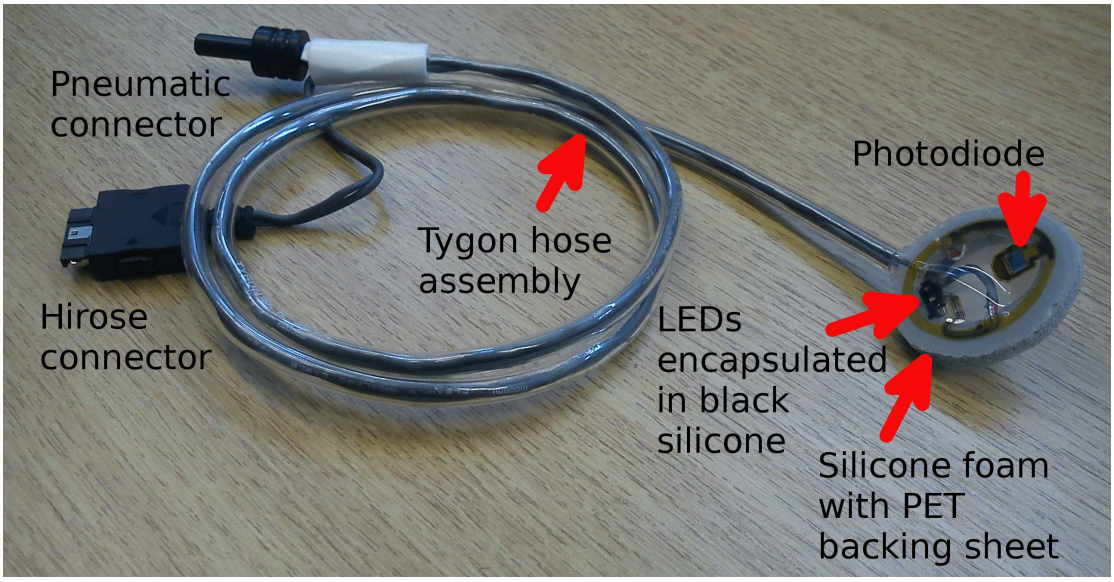


Figure 9.12: Silicone rubber membrane based pneumatic capillary refill sensor.

Figure 9.12 shows a labelled photograph of the assembled membrane sensor. The hose assembly consisted of bio-compatible Tygon tube (Saint-Gobain part number S-50-HL with an inside diameter of 3.8mm and outside diameter 4.8mm) with a three core shielded (Mogami Cable "Ultraflex" 2879-08, used to drive the three LEDs with 660nm, 810nm, and 950nm peak emission wavelengths), and two core shielded (Mogami Cable "Ultraflex" 2490-08 used to interface with the photodiode and a $10k\Omega$ thermistor) cable passed through its centre. The gap between the cable and inside face of the hose was found to be sufficient to pass the air supply through to the sensor.

A thermistor (Murata NCP15XH103F03RC) was mounted on the flexible polyamide PCB, placing it in more effective thermal contact with the skin surface than was the case in the prototype sensor. This was found to result in a significant decrease in

the response time, and with the insulation provided by the silicone foam rubber, the temperature data provided by the thermistor should be a more accurate measure of skin surface temperature, although this was not exhaustively tested (refer to figure 9.14 for example temperature data recorded from the sensor). Although the materials used to construct the "proof of concept" sensor shown in figure 9.12 are not bio-compatibility certified, silicone rubber membrane film and silicone foams with ISO10993 certification are available, albeit harder to source in small quantities for prototyping use. The non-biocompatible silicones were used only due to difficulties sourcing suitable materials, but for a production sensor this should not be a problem.

The sensor's source to detector separation was approximately 15mm. This was found to be too large for a useful signal to be obtained. It was also found that there was excessive ambient light leakage through the white plastic PET (Polyethylene terephthalate) sheet used as the upper face of the sensor. For these reasons a second membrane sensor was assembled with a shorter source to detector separation of 6mm and black PET sheet (125μm thick) used for the upper face (figure 9.13).



Figure 9.13: Black silicone sensor, 6mm LED to photodiode separation

A more serious problem with the membrane sensor design was the difficulty of bonding silicone rubber to the plasticised PVC hose assembly. Although a strong bond could be made using RTV silicone, the bond strength was found to decrease over time, with complete failure after approximately 2 weeks at room temperature. This is believed to be due to plasticizer migration through to PVC to the bond region; accumulating plasticizer on the PVC surface then leads to failure of the bond to the silicone rubber. One possible solution to this problem would be to use an interface material consisting of either a plasticiser free plastic (such as a rigid, high surface energy polymer e.g. Nylon 6,6) or a metal shroud glued around the end of the PVC hose.

Durability posed another potential problem with the silicone membrane sensor; in order to apply adequate pressure to the skin surface it was found necessary to use extremely thin 150μm silicone rubber film for the sensor bottom. To avoid coupling artefact problems, the bottom face of the sensor was coated over its entire surface with Dow Corning MD7-4502 medical device silicone skin adhesive, silicone adhesive being necessary to ensure adhesion to the rubber film. The strong adhesion of the adhesive to the skin

caused significant distortion of the membrane as the sensor was peeled off the skin surface after use, and was found to detach the LED cluster and the photodiode from the membrane of the first sensor after only a single use of the device on the forearm of a healthy volunteer.

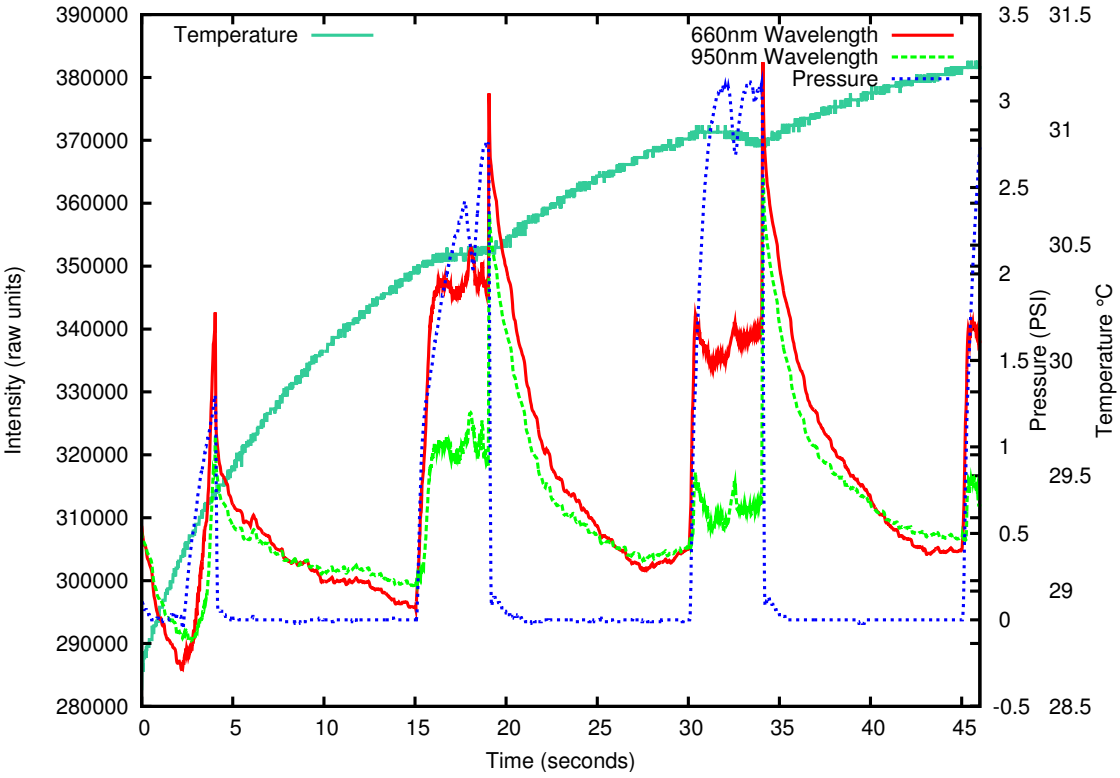


Figure 9.14: Raw intensity data from a capillary refill test using the black silicone sensor, applied pressure and temperature are also plotted.

Figure 9.14 shows capillary refills successfully recorded using the second (figure 9.13) silicone rubber sensor. In order to avoid the risk of sensor damage on removal from the volunteer's skin, no adhesive was used, and instead a 1Kg mass was placed on top of the sensor to hold it in position. It would appear that this led to several problems with the recorded data: the raw intensity increases immediately following the pressure release, an effect that is thought to be due to either the LEDs or the photodiode area (or both) of the membrane losing contact with the skin surface when the air pressure is applied; and the temperature drops slightly during each pressure application period, again suggesting that the sensor face is losing contact with the skin surface, allowing the thermistor to cool slightly. Varying source to detector separation as the membrane distorts is another possible explanation for the observed intensity artefacts. Zysset et al describe separation induced artefact problems with a textile based NIRS system [234]. However, the decreasing temperature during the pressure application periods and the observation that the sensor tended to start to leave the skin surface during pressure application

would suggest that variations in optical coupling explain the majority of the observed intensity offset. If varying source to detector separation is indeed a significant source of artefact, it may be possible to improve the design with polyamide PCB material between the LEDs and the photodiode. This would help to maintain a consistent separation due to the high modulus of polyamide film. It should also be noted that the observed artefact effect is not necessarily a serious impediment to successful data processing, as it would appear that the separation/coupling effect does not give a variation in intensity during or after the refill period itself (i.e. after the pressure is released).

It should be noted that in figure 9.13 three LEDs can be seen, with an 810nm LED at the far left followed by 950nm and 660nm moving to the right of the cluster. Unfortunately an error during testing led to the 810nm LED being damaged, so only dual wavelength data could be recorded in the volunteer test.

It can be seen that, compared to the transcutaneous oxygen sensor based prototype (data plotted in figure 9.11), there is an increase in both the relative amplitude of the capillary refill intensity change and the time taken for the intensity to return to baseline or refill time. The large amplitude response can be explained by a combination of the 6mm path length, which is longer than the path length of under 5mm in the prototype trans-cutaneous attachment ring sensor, and the reduced light leakage through the sensor, meaning that almost all of the detected light had passed through the skin. The longer refill time may be as a consequence of the large mass placed over the sensor restricting blood flow, or simply as a result of changes in the vascular geometry between measurement sites (the silicone sensor was tested in the centre of the palm due to its size, whereas the trans-cutaneous attachment ring was tested on the Thenar Eminence).

The durability problems highlighted the difficulties inherent in working with silicone rubber based materials; the surface energy of silicone rubber is extremely low (the surface energy of PDMS silicone rubber is $18mJm^{-2}$, compared to for example $45mJm^{-2}$ for polyethylene terephthalate [235]), and it is only possible to bond silicone using other silicone rubber materials (RTV silicone and pressure sensitive silicone resin based adhesives were used in the capillary refill sensors). Even when bonding two silicone rubbers together, the bond strength is weak, leading to the durability problems that were experienced with the sensors. In addition, it is difficult to achieve a strong and long lasting bond between silicone rubber and other polymer materials. For these reasons it was judged that a silicone rubber based sensor would not be the most appropriate solution for a production capillary refill sensor.

Bio-medically approved polyurethane films are widely available, for example from Smith and Nephew PLC in the UK. These are also available in a water vapour permeable micro-

porous form, preventing the accumulation of sweat beneath the device in a long term use scenario. Silicone rubber has a much higher water vapour permeability than almost all other solid polymers, but only a thin layer of silicone has sufficient permeability to allow sweat to dissipate effectively. If sweat is allowed to accumulate beneath an adhered medical device, adhesive strength decreases due to separation of the top layer of skin cells in the epidermis [236].

An additional advantage of polyurethane over silicone rubber is ease of manufacturing, since polyurethane is available as a thermoplastic elastomer, meaning it can be put through manufacturing processes such as vacuum forming. Silicone rubbers are only available in thermosetting form, either as a two component product or a room temperature vulcanising (RTV) silicone cured by atmospheric moisture. Cured silicone rubber is extremely resistant to high temperatures, but begins to irreversibly thermally decompose above 300°C, meaning it must either be moulded into shape before it cures or cut to shape; a process that is limited by the low modulus of the material, meaning cutting is the only suitable method of machining silicone film components.

Based on the experience gained from the two silicone rubber sensors a polyurethane film based sensor was designed with mass manufacturability in mind.

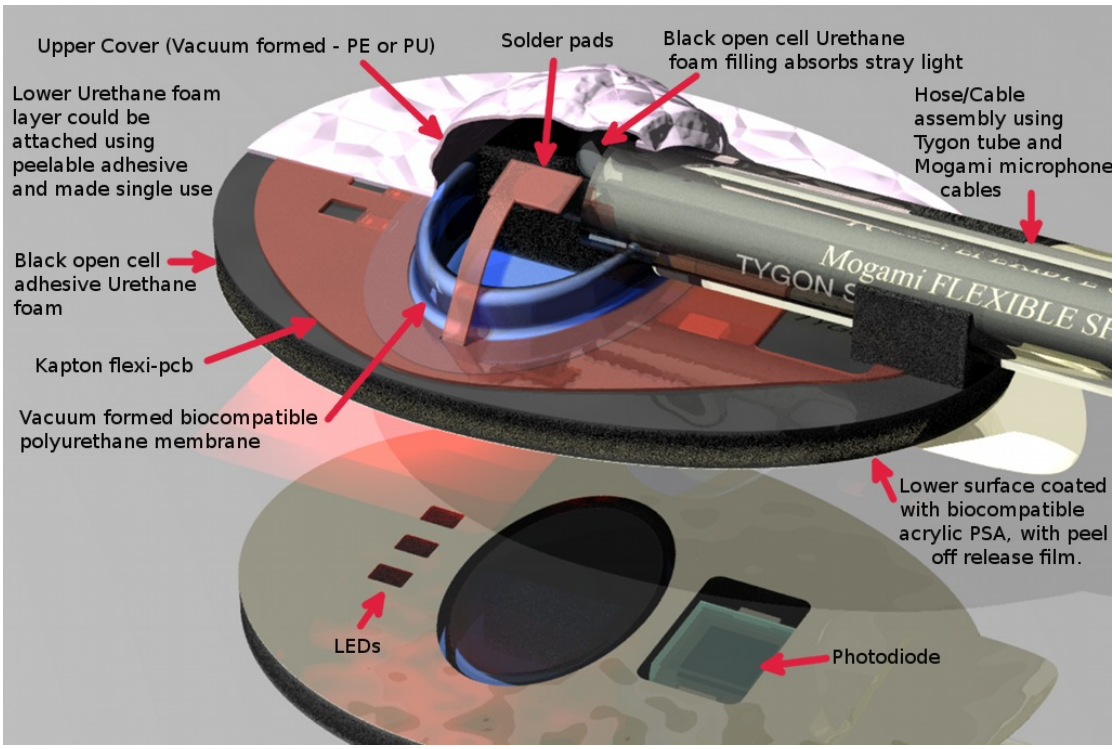


Figure 9.15: Rendering of polyurethane sensor design proposal.

Figure 9.15 shows a CAD rendering of a cross sectional view of the proposed sensor. The goal of this design was to allow for vacuum forming and die punching of all the sensor components. The majority of the components could then be assembled using adhesives,

for example flexible cyanoacrylate derivatives such as butyl-cyanoacrylate (e.g. Loctite 4851), which are available as certified bio-compatible products. Alternatively, medical grade urethane or light cured adhesive might be used. In the design proposal, the LEDs and photodiode have been placed on either side of the pressure applicator in an attempt to avoid the variable optical coupling effects seen in the silicone rubber sensor. By maintaining an air gap between the optodes and the skin surface it is hoped that this sensor design would see an approximately constant intensity loss at the skin surface, and by using black urethane foam light piping through the sensor itself could be avoided, leading to a stable signal throughout the pressure application and release cycle.

It may also be feasible to use this proposed design in a reusable form for hospital trials; by using a peel-able adhesive to coat the bottom surface of the polyamide flexible PCB, it would then be possible for the lower conformable polyurethane foam skin contact layer to be replaced after each use, necessary since this foam layer would be likely to sustain damage during sensor removal and would be extremely hard to clean effectively ready for use on the next patient or volunteer.

Based on feedback on the first polyurethane sensor design proposal from clinicians, a second, revised sensor design was prepared. The LEDs and photodiode were incorporated into the pressure application diaphragm, increasing the durability of the device at the expense of an increased risk of signal artefacts due to variable optical coupling with the skin surface. A single sheet of vacuum formed polyurethane was designed to form the sides and bottom surface of the sensor (i.e. the surface in contact with the skin), with the top surface being formed from a sheet of 150μm PET, followed by a bio-compatible layer of thin, adhesive coated polyurethane film. These changes resulted in a considerably smaller sensor of 23mm outer diameter, and 5mm total height.

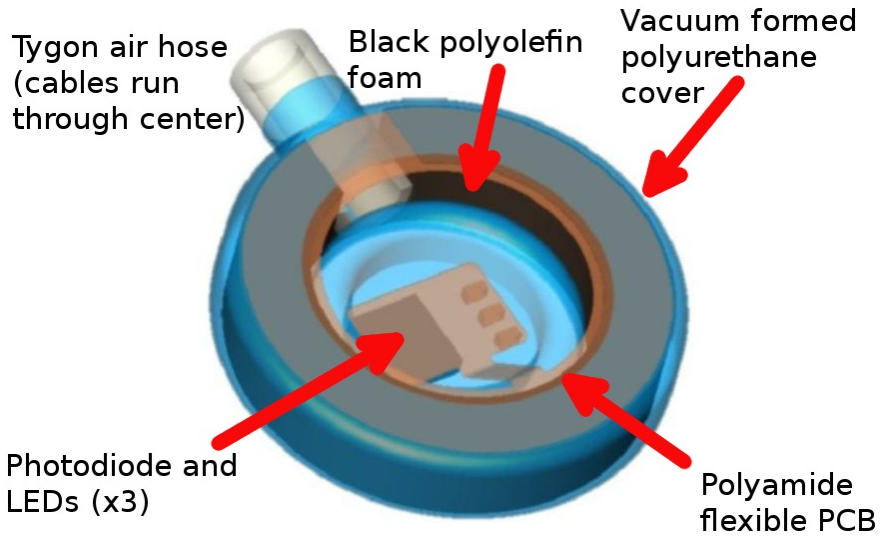


Figure 9.16: Rendering of second polyurethane sensor design proposal.

Figure 9.16 shows a CAD rendering of the second polyurethane sensor design with the PET top sheet removed for clarity. A 1.5mm deep corrugation is formed into the diaphragm, allowing it to easily distort when air pressure is applied to the device. The cables passing down the centre of the air hose are also not rendered. In the finished device these would be soldered to the polyamide flexible PCB.

As the outer surface of the sensor is formed from only two pieces of polyurethane in this new design, it is possible to clean the sensor for limited reuse. This would be a significant advantage if the device was to be used in a small or medium scale study, as only a small number of sensors would be required, making hand assembly practical.



Figure 9.17: The vacuum formed polyurethane sensor cover. The filament deposition moulded ABS mold can be seen, along with an MDF adaptor sheet.

A single prototype polyurethane based sensor of the second design was assembled in order to assess the practicality of the design. A breathable, bio-compatible (ISO10993) polyurethane sheet product manufactured by Smith and Nephew Plc. (EU95, a 95μm grade sheet) was selected for use as the skin facing side of the sensor outer cover. This material has the advantage of existing widespread use in medical plasters and bandages, and a selection of acrylic skin adhesives which can be coated onto the sheet are available from the same manufacturer. In volume production, adhesive coating could be applied to the bottom face of the sensor, followed by a release layer (e.g. PTFE coated paper), which would be peeled off the sensor before it was applied at the measurement site. Unfortunately, removal and reapplication of the adhesive for sensor reuse is unlikely to be practical, so securing the sensor from above with tape is likely to be necessary if the sensor is to be easily reused.

A fused deposition molding machine (Statsys Inc.) was used to make the a positive mold for vacuum forming the outer cover. The vacuum formed sheet can be seen in figure 9.17. Considerable difficulties were encountered with wrinkling during forming, as can be seen in the figure; the pictured formed sheet had to be discarded for this reason. It was found that heating the polyurethane close to its melting point reduced the risk of

creasing. A non contact infrared thermometer was used to check the film temperature during heating, and it was found that the most reliable forming temperature was 180°C. It is likely that forming reliability could be significantly improved by optimising the shape of the free sheet around the mold to find the best area draw ratio.

The same 4.8mm outer diameter Tygon PVC tube, with Mogami "ultraflex" cables threaded down its centre, was used as in the previous sensors. A closed cell polyolefin foam (5mm thick Alveolit from Sekisui-Alveo AG.) was selected for use around the outer circumference of the sensor, based on its relatively high stiffness. With an internal air pressure of 3PSI inside the sensor, the foam is under a tensile load of approximately 1.5PSI. The grade of foam was chosen such that elongation under 1.5PSI was less than 25%. The Alveolit foam can be accurately cut with a 120W CO_2 laser cutter, a technique that may also be applicable to volume sensor production. Die cutting foam parts with a high thickness to width ratio whilst maintaining parallel edges is challenging.

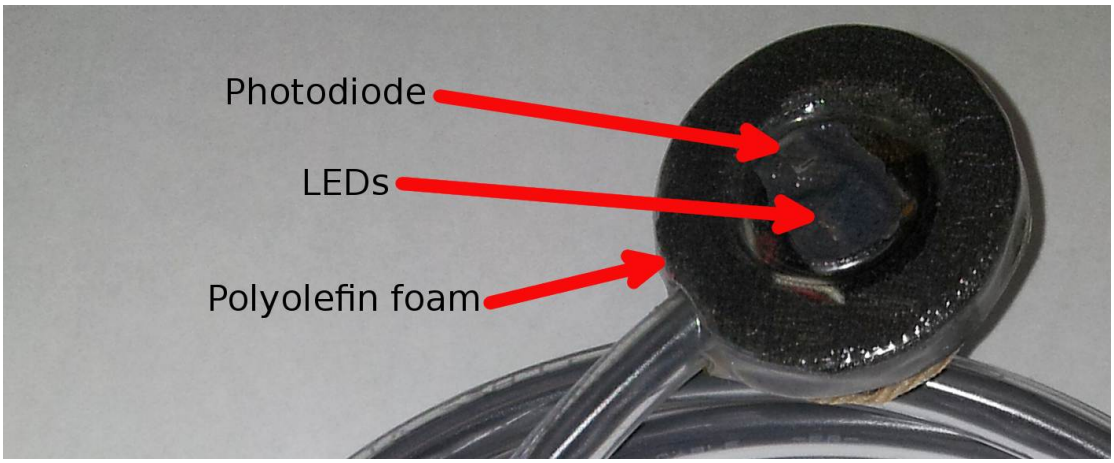


Figure 9.18: The completed polyurethane sensor prototype.

Figure 9.18 shows a view of the assembled prototype sensor from the bottom or skin facing side. An opaque epoxy resin encapsulates the LEDs and photodiode (a Vishay BPW34S), to prevent direct optical coupling and secure them to the face of the flexible diaphragm. The flexible PCB is not visible in the figure, this uses a 125μm polyamide substrate with 1oz/ft^2 copper coating.

Unfortunately it was found that the vacuum formed polyurethane outer coating did not have a sufficiently wrinkle free surface for the PET back sheet to form an airtight seal. Further refinement of the molding process is clearly required. Use of a negative forming mold is one possibility, and a top mold may also be required.

9.3 Continuous wave NIRS data logger design for low volume production

In order to both demonstrate the applicability of the continuous wave NIRS system for manufacturing scale up and also enable a large number of trial recordings to be made from both healthy volunteers and potentially also patients, it was decided to design a second revision of the data logger electronics. The changes to the new design were based on the experience gained operating the first revision of the hardware in experiments on healthy adult volunteers.

9.3.1 Electronic design

From figure 9.3 it can be seen that the SMPS takes up a large proportion of the board area. The high quality inductors and capacitors that were required for acceptable performance of the SEPIC converter also added considerably to the Bill Of Materials (BOM) cost. With these facts in mind it was decided to investigate the gain in total data logging time that resulted from use of a switching supply. The prototype data logger was left running with a manometer connected to the air supply connector, and one of the trans-cutaneous attachment ring based sensors connected to the Hirose sensor interface. The battery discharge curve under typical operating conditions could then be logged.

From figure 9.19 it can seen that there is a "knee" in the cell voltage at 3.65V, with less than 3% of the cell capacity remaining once the voltage reaches this level. This is 350mV above the 3.3V voltage used for the analogue and digital supply rails of the logger electronics, so there is more than sufficient voltage headroom to allow the device to function effectively with linear regulators connected directly to the cell voltage.

If low dropout regulators with a dropout voltage of 100mV are used, then it can be seen from the graph that there will be a negligible decrease in the fraction of cell capacity used by the logger. At low cell voltage of around 3.6V the linear regulator loss will be approximately equal to the power loss in the SEPIC, however, when the cell is at a higher charge level the SEPIC converter's switching efficiency of approximately 90% means that less current will be drawn from the cell. Extrapolating from the curve it can thus be seen that use of the SEPIC converter results in an increase in data logger operating time of less than 10% at typical operating temperatures. Thus it was decided to remove the SEPIC converter from the design and adopt a power supply scheme with multiple switched low dropout linear regulators for the multiple rails driven directly for the lithium polymer cell, i.e. the schematic from figure 9.2 with the SEPIC removed and the 3.5V SEPIC rail replaced by the cell rail.

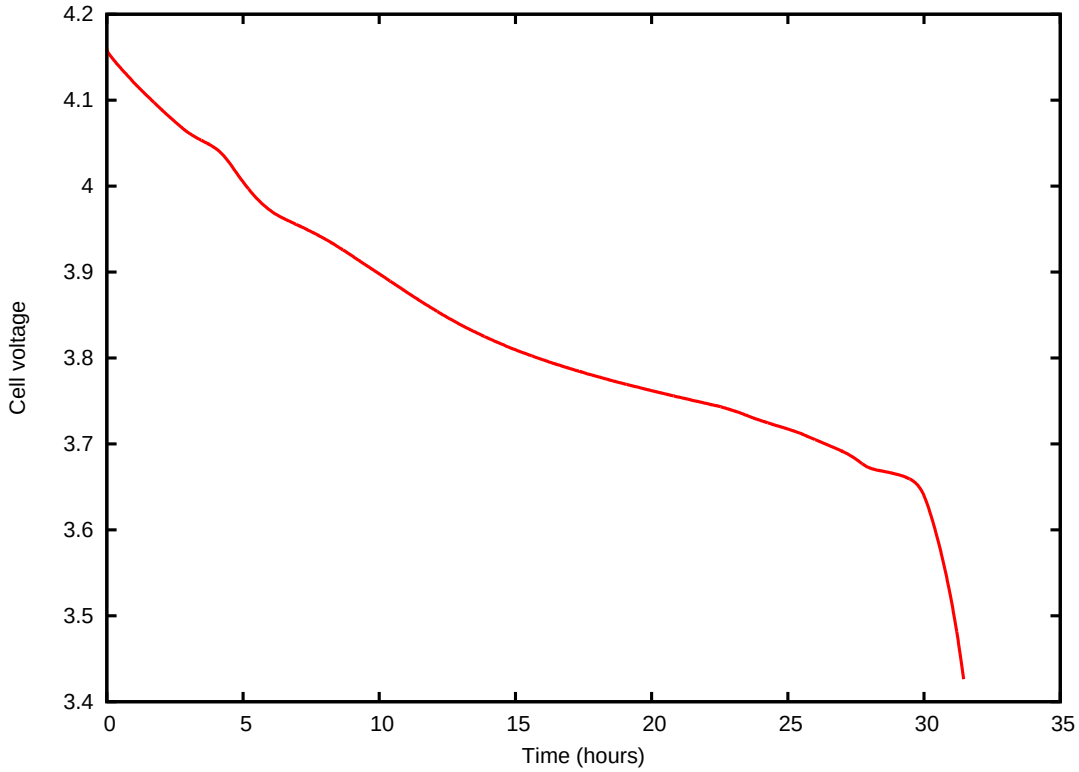


Figure 9.19: Discharge curve for the prototype data logger. A 1.3Ah lithium polymer cell (BAK LP-503759-1S-3) was fitted and the device was maintained at room temperature throughout (20 to 25°C).

The pressure sensing and battery charging circuits were also redesigned in order to make cost savings and reduce board area on the new PCB revision. Based on the tear-down of the Omron M2 automated blood pressure meter, the Honeywell sensor was swapped for a considerably cheaper and more sensitive Omron sensor, the 2SMPP, a single port gauge sensor designed specifically for blood pressure monitoring and having an operating range of 0 to 5.4PSI, well matched to the requirements for the data logger. The 2SMPP sensor requires a constant bridge current of 100μA to maintain accuracy over temperature, making sensor interfacing slightly more complex than with the earlier Honeywell sensor. To this end, an adjustable voltage regulator was used as a simple constant current supply for the sensor by powering the strain gauge bridge from the voltage regulator output and placing a 120$k\Omega$ resistor between the sensor bridge ground and the board ground, ensuring a constant 100μA current through the sensor when used with an adjustable regulator with a 1.2V internal reference voltage.

The MAX1811 lithium ion charging IC was replaced with the smaller and lower cost BQ24090 from Texas Instruments. The BQ24090 has the additional advantage of built in temperature compensation using an external thermistor (although this is relatively rudimentary). Due to the constant current drive technique used for the thermistor sense

pin on the BQ24090 it was found that board temperature measurements could be made by connecting ADC2 on the STM32 micro-controller to the "TS" pin on the charger. The STM32F1 series micro-controller has an internal die temperature sensor, but this is poorly calibrated and has a relatively non-linear response. The internal sensor can also only be used with ADC1, which was the only ADC unit capable of being used with the NIRS due to its superior DMA capabilities compared to ADC2, and so STM32 die temperature measurements could only be made if NIRS spectroscopy was temporarily disabled.

The pin header and separate analogue front end sub module was also replaced with an on-board front end, using a single SOT23-5 operational amplifier (MicoChip Inc MCP6L91), with the single op-amp front end design from Chapter 7 being used.

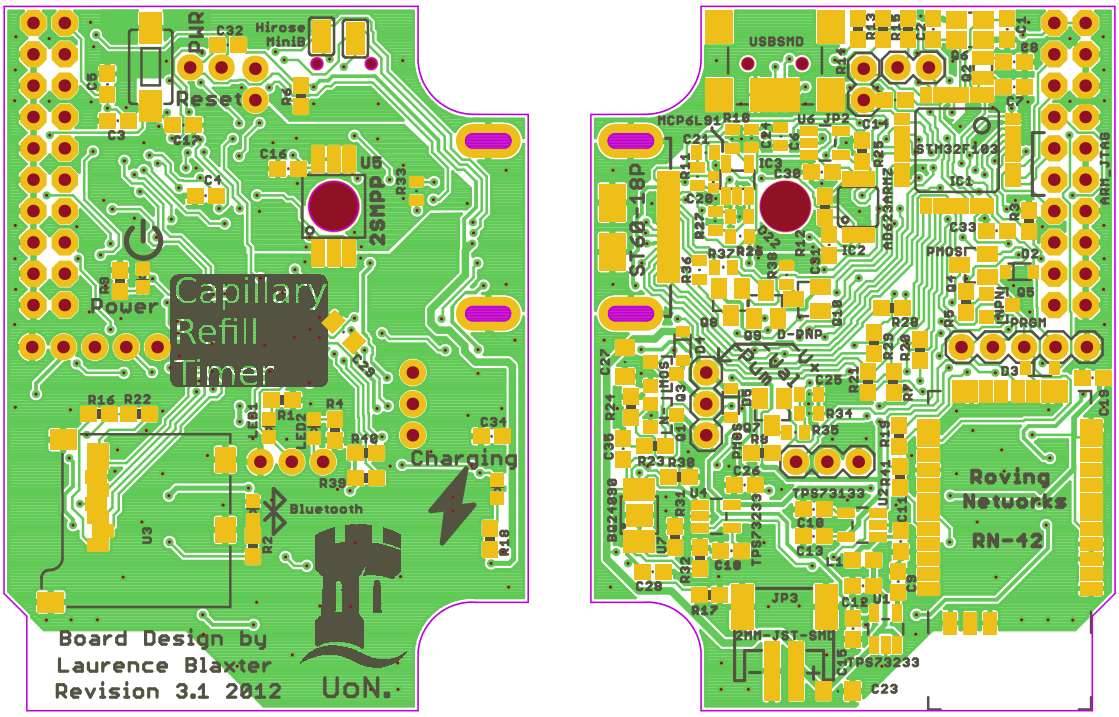


Figure 9.20: The PCB design for the second data logger revision, designed for low volume production.

Figure 9.20 shows a view of the bottom (left hand side) and top faces of the two-sided PCB designed for the second data logger revision. As in the previous design, the "bottom" face of the PCB was designed to face upwards inside the translucent polycarbonate enclosure enabling the indicator LED labels to be read. The layout was made considerably more compact, and together with the simplifications to the schematic this reduced the overall board size to 42x57mm, a reduction in board area by a factor of 2. By integrating the Bluetooth radio onto the main board, moving to a 4 layered

design and removing the USB connector and 20 pin 0.1 inch JTAG connector in favour of a more compact Cortex debug connector it would likely be possible to reduce the board area by another factor of between two and three. However, as the overall size of the base unit is limited by the size of the pneumatic components (diaphragm pump and solenoid valve) and the 1.3Ah lithium polymer cell this was not pursued.

9.3.2 DSP bandpass filtering

When the data logger was first used in a clinical environment (as part of the healthy volunteer study discussed in Chapter 10), the NIRS was found to be vulnerable to optical interference from modulated light sources such as LCD monitors and some fluorescent lights. The frequency division scheme (section 7.2.3) demodulates via a rectangular windowed DFT, leading to significant sidebands. Although unproven, it was suspected that modulated ambient light "leaked" into the demodulated signal via these sidebands.

A potential solution to this problem would be to multiply the ADC data by a windowing function such as Blackman-Harris. However, this would result in a decrease in signal to noise ratio, and the relatively large windowing function lookup table would need to be stored in the flash memory of the STM32 micro-controller due to the insufficiently large RAM. Reading from flash would cause bus contention with the processor core, which runs directly from the flash memory, leading to a serious decrease in processor performance as the core would regularly enter wait states.

For these reasons, an alternative strategy was employed; an IIR bandpass filter was placed between the ADC and the demodulator, using a second order direct form 1 biquad implementation written in Thumb 2™ assembler. It was found that a filter centred on approximately 11700Hz (just below the 11905Hz central modulation frequency) with a 3dB bandwidth of slightly over 1kHz could be implemented using coefficients of the form $1 - 2^n$, where n is an integer, allowing simple bit shift operations to be used

$$y_n = x_n - x_{n-2} + 2(1 - 2^{-7})y_{n-1} - (1 - 2^{-7})y_{n-2}. \tag{9.1}$$

The filter difference equation (9.1) was implemented using a SIMD (Single Instruction Multiple Data) scheme, processing two samples simultaneously using the upper and lower 16bit words of a 32bit register. This filter used 18 core clock cycles per ADC sample, giving a total processor overhead of 16Mhz, well within the spare capacity of the 72Mhz STM32.

A reduction in the incidence of interference was seen once this filter was used on the data loggers in the volunteer study.

9.3.3 Cost analysis

To reduce the PCB manufacturing costs the PCB design in figure 9.20 was panellized into a 3 by 4 array. This was sized to allow the boards to be populated using conveyor belt pick-and-place / linear reflow processing, potentially offering a significant reduction in assembly costs in large volume production. For the trial production run, a single panel of 12 units was populated by hand using a laser cut stainless steel stencil to apply solder paste followed by soldering using a reflow oven.

Including the cost of the pneumatic components, the total cost of a single base unit when ordering parts in units of 10 came to £162 ($260) excluding taxes. This falls to $180 if the pneumatic components (pump and valve) are excluded. In larger quantities this cost is likely to decrease significantly; applying typical scaling laws for manufacturing scale-up in the electronics industry, the per unit cost would be likely to reach around $80 in 10,000 unit quantities. The cost of the pneumatic components is difficult to estimate, and it would appear that the Parker-Hannafin prices are excessive compared to other manufacturers, for example the Omron M2 blood pressure monitor retails for around $50, and the pneumatic system includes a diaphragm pump and solenoid valve from a Chinese manufacturing group. Thus it would seem reasonable to conclude that the pump and valve could likely be sourced for less than $20 in large scale production quantities, bringing the total per unit cost to under $100.

This is extremely cheap compared to typical medical device costs, for example pulse oximetry base units typically have list prices in the region of $2,000. Thus it can be seen that the data logger base unit design is suitably low cost to not only enter widespread use with an acceptable allowance for inclusion of development/regulatory costs and profit margin, but also for novel use cases to emerge, for example in the developing world.

9.3.4 Enclosure design and assembly

The same translucent polycarbonate enclosure (Hammond Inc. 1591BTBU) as used in the prototype was again used for the trial manufacturing run of data loggers. For large volume production it would be likely to be more cost effective to redesign the PCB to fit in a custom built enclosure, but manufacture of injection moulding tools for small production runs is not cost effective. As well as dust and splash proofing, the Hammond enclosure also provided the 6kV ESD protection sufficient to meet the EN61000-4-2 requirement.

In order to meet the EN980/ISO15223:2012 standards for device labelling, a durable printed foil based labelling system was applied to the bottom face of the enclosure.

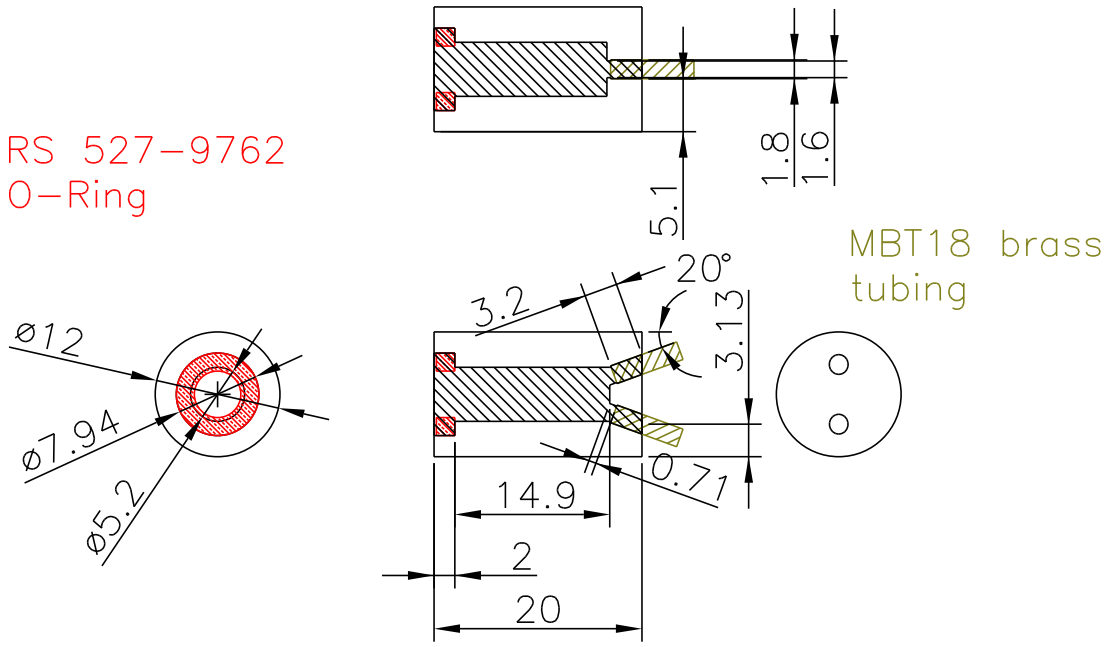


Figure 9.21: The custom pneumatic connector design for the second data logger revision, designed for low volume production.

The multifunction on/off/control button was also labelled with the IEC 5009 symbol to appropriately indicate its function.

To reduce the cost of the base unit, the Omron M2 pneumatic connector seen in figure 9.5 was replaced by a custom made but compatible design machined from Nylon 6,6 (figure 9.21). A silicone rubber o-ring was used as an air seal, with two 1.8mm diameter by 1cm long lengths of brass tube being fixed into the inside face of the connector for use as hose nipples. The Nylon connector assembly was then glued to the inside face of the polycarbonate enclosure using a cyanoacrylate adhesive.

Self adhesive foam (sourced from 3M) was used to secure the Lithium polymer cell, diaphragm pump, and solenoid value within the enclosure. The off-the-shelf button assembly used for the on/off/control button incorporated a retaining clip mechanism, so no adhesive was required to secure it to the enclosure. Rubberised adhesive feet were placed on the bottom of the enclosure to minimize the risk of it slipping off desks or bedside tables etc. in use. Figure 9.22 shows an external view of one of the completed data logger units.

The slots and holes in the Hammond enclosures were milled using a three axis CNC machine operating from a CAD design. The electrical interface connector (used for the

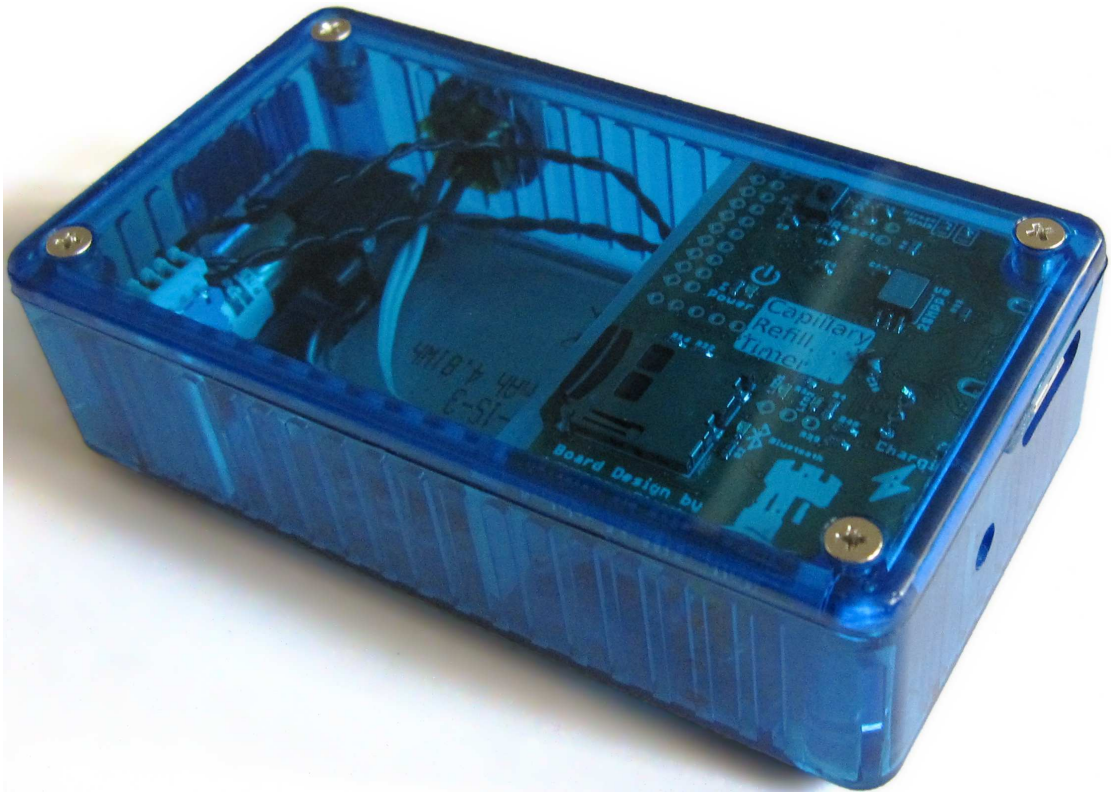


Figure 9.22: An assembled data logger unit - final hardware revision.

USB interface and sensor connection) is visible as the slot in the front (right hand side) face of the enclosure, and the pneumatic supply connector as the hole to the lower left hand side of the interface socket.

A total of eight complete data loggers were assembled, allowing relatively large scale clinical trials to be conducted into a range of CW-NIRS based techniques (Capillary refill, venous occlusion, or multi wavelength photo-plethysmography). Three of the units were used in the PaedCEW study, which is discussed in further detail together with preliminary results in Chapter 10.

It may also prove possible to conduct more conventional measurements, such as blood pressure or pulse oximetry, due to the ability to develop custom firmware for the devices.

Chapter 10

Experimental optical perfusion measurement in healthy volunteers

In order to further explore the potential utility of the optical perfusion monitoring techniques identified in Chapter 7, a series of studies were conducted on healthy adult and child volunteers.

Data from a healthy volunteer study assessing the automated capillary refill device developed in Chapter 9 (the "PaedCEW" study) is presented in this chapter. In this study, laser Doppler, Masimo PI perfusion monitoring, and automated capillary refill timing (based on a simple intensity threshold model) were conducted simultaneously on 17 child (aged under 16) and 20 adult volunteers during exposure to a 5°C chamber, allowing the relative performance of these three candidate techniques to be assessed.

As the sensor hardware used in the PaedCEW study was not sufficiently well characterised for NIRS techniques to be applicable, a second small scale study of adult volunteers was conducted, using NIRS to further understand the blood flow mechanisms responsible for capillary bed refilling in the capillary refill test.

Two of the optical techniques discussed in Chapters 6 and 7; the NIRS combined with venous occlusion technique, as well as the pulse oximetry derived "Perfusion Index" (or PI) perfusion metric, have seen considerable use in published perfusion research. It is for these reasons, in addition to the doubts concerning the true nature of the capillary refill measurement (i.e. how closely it is related to perfusion), that it was decided to devise an experimental protocol for this second study that would attempt to benchmark these three optical techniques against one another.

10.1 Experimental protocols

10.1.1 Proposed protocol evaluating capillary refill, venous occlusion, and Masimo PI

A simple benchmarking process could consist of three simultaneous baseline measurements recorded from each volunteer using the three different techniques. However, the PI metric output from Masimo pulse oximeters cannot be converted directly to a volumetric perfusion, and interpreting NIRS recorded from the automated pressure application device described in Chapter 8 proved to be challenging. So for these reasons, a method for inducing a change in peripheral perfusion of the experimental subjects would be desirable, in order that the response of these three different techniques could be compared.

Ideally this method would be non invasive and have minimal potential for injury to the subject. Two such potential methods were identified:

1. Changing the environmental temperature of either the entire limb or the hand or foot under study, for example by use of a water bath. Such changes lead to vasodilation (on heating) and vasoconstriction (cooling). This alters the resistance of the vascular bed to blood flow, and in turn leads to a change in perfusion.

   A water bath or chamber has the advantage of safety and repeatability, as the temperature can be set and measured accurately. However, the response of the sympathetic nervous system in regulating blood flow in response to temperature may differ considerably between subjects, so although the bath temperature could be well controlled, there would be uncertainty as to the true change in perfusion.

   It can also be seen that there is the potential for a misleading negative correlation between environmental temperature (and by implication perfusion) and capillary refill time, as a decrease in capillary refilling time with increasing temperature could arise due to both the increase in inflowing blood from the arterial supply (i.e. the true perfusion), and due to increased mobility of blood within the vascular bed due to vasodilation, with this allowing blood to flow into the pressure affected and blanched region from surrounding regions more quickly.

2. Use of an off the shelf, electronic blood pressure measurement cuff to cut off blood supply to the arm, and thus changing perfusion from baseline to zero.

   This has potential hazards for the subject, as an inflated cuff applied for an excessive time might lead to tissue damage in the arm, and if overinflated the cuff may

cause bruising or more serious injuries. However, commercially available blood pressure measurement devices incorporate safeguards to minimize these risks.

Such an electronic blood pressure cuff would be capable of applying pressures larger than systolic blood pressure only for short periods (of order a minute). Although this is sufficiently long for a capillary refill measurement to be taken, measurement of PI might be inaccurate.

Venous occlusion could be made using of a second cuff placed lower down the arm and inflated to a lower pressure (50mmHg has been used in most studies), or by inflating the pressure cuff used for arterial occlusion to a lower pressure. Successful cutoff of the arterial flow by the cuff could be tested using a NIRS device placed on the hand to check for an absence of both pulsatile fluctuations and any trend in total tissue haemoglobin with time.

It was decided to draw up an experimental protocol making use of both of these techniques, as it was judged that both arterial occlusion and water bath induced environmental temperature changes can be carried out with minimal risk or discomfort to the experimental subject, and that both techniques may provide a useful insight into perfusion measurement.

Simultaneous measurement of capillary refill with combined venous occlusion and NIRS would be technically feasible using the hardware discussed in Chapter 8, with the trend in baseline intensity level between refills being utilised to extrapolate change in blood volume. However, it is unclear whether the venous blood pressure that develops during a venous occlusion leads to a biassing of the capillary refilling, for example, intuitively it might be expected that as the venous pressure rises there will be an increased rate of capillary bed backfilling from the venous blood compartment.

For this reason it was decided to conduct capillary refill and venous occlusion measurements at separate times. Using an off the shelf blood pressure cuff allows pressure to be ramped up over a relatively short time (less than twenty seconds) for arterial occlusion. This means that the time lag between venous and arterial blood flow cut off times will be sufficiently short to avoid significant blood pooling within the arm (published studies indicate that pooling takes at least a minute, e.g. Zysset et al [234]), and thus it should prove feasible to combine arterial occlusion with capillary refill without significant biassing of the refill data.

Figure 10.1 shows a schematic diagram of the proposed protocol. The arterial occlusion is indicated by the red hatching before and after the water bath immersion. As the Masimo pulse oximetry finger probes can be water proofed, one can be applied continually to the

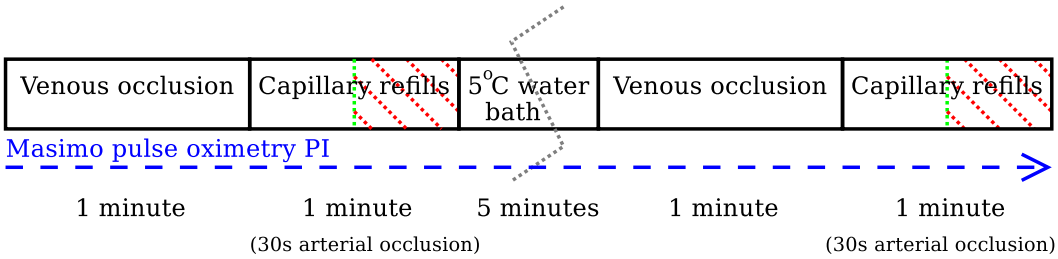


Figure 10.1: Proposed perfusion state modification and measurement experiment.

finger throughout the entire experiment, and the PI perfusion metric recorded at a 1Hz sampling rate via the RS232 serial port on the machine. The Omron M2 electronic blood pressure meter (subject of the Chapter 9 teardown) allows pressure to be set manually, so a single blood pressure cuff applied to the upper arm throughout the experiment and inflated via an Omron M2 unit can be used for both venous occlusion (in manual mode with a pressure set to 50mmHg) and arterial occlusion (in normal mode). Meanwhile, the automated pressure application and NIRS device discussed in Chapter 8 allows NIRS to take place during both baseline conditions and also application and release of pressure (i.e. capillary refill) to the hand or forearm.

A 5°C water bath temperature can easily be achieved using a domestic refrigerator, and is not an unduly unpleasant temperature for 5 minute immersion of the hand.

To summarise, the experimental plan shown in figure 10.1 should allow the effect of low environmental temperature on capillary refill, total haemoglobin perfusion derived via venous occlusion and NIRS, and the pulse oximeter derived PI metric to be measured in a consistent manner. Use of arterial occlusion combined with capillary refill both before and after the water bath immersion will allow the effect of environmental temperature on capillary refill in the unperfused state to be observed. Absence of a baseline shift from the spectrometer during the arterial occlusion, combined with a loss of pulsatile signal from the pulse oximeter indicating successful occlusion of the arterial supply to the arm.

Any change in the capillary refill between the two arterial occlusions will indicate effects that are purely due to the temperature induced vasoconstriction. Thus it can be seen that the protocol described in figure 10.1 allows an insight into the true nature of the capillary refill, and potentially a better understanding of the relative contributions of perfusion blood and blood inflowing from surrounding tissue in the capillary refill test.

10.1.2 Use of data from "PaedCEW" research project

The "PaedCEW" (Paediatric Continuous Early Warning) project at the University of Nottingham involves research into automation of the capillary refill test. This work was funded by an MRC DPFS grant and it allowed further opportunities to perform mea-

surements on patients in a controlled environment. As part of this study, an automated pneumatic capillary refill device using the base unit described in Chapter 9, together with a custom made pneumatic sensor developed independently by Dr David Morris (Appendix A) was tested on a group of healthy child and adult volunteers. This study was carried out with the assistance of clinical research fellows Dr Caroline Henry and Dr Sarah Hill at the School of Medicine, University of Nottingham, and was conducted in accordance with the University of Nottingham code of research conduct and research ethics [237]. The pneumatic sensor is described in the appendix in more detail. Briefly it used a three wavelength reflection mode photo-plethysmography (PPG) device mounted inside a rigid enclosure, with a size 1 (3 by 6 cm) blood pressure cuff secured between the back of the PPG sensor and the enclosure, allowing the pneumatic supply from the data logger unit to be used to vary the application pressure.

The PaedCEW experimental protocol involved a 20 minute baseline recording followed by a 20 minute low temperature exposure period, during which the forearm was placed inside a custom made chamber at a temperature of 4°C. The chamber is shown in figure 10.2. Two iris type insulated foam portholes can be seen in the left on the chamber, together with two circular insulated portholes in the door. Due to the relatively good seal formed by the insulating foam, the experiment participants were instructed to place their arms into the chamber using the side portholes, with the front portholes used for cabling and hose access to the sensors.

A custom made control unit (sourced from the medical engineering unit at the Queens medical centre, Nottingham) installed on the top of the chamber unit allows the setpoint temperature to be controlled, and the current internal temperature can be displayed via a seven segment LCD. This control unit can be seen in the figure on the top left of the chamber.

Four separate devices were used throughout the experiment; a Moor instruments (MoorVMS) point source laser Doppler device, a Masimo Radical 7 pulse oximeter, a pneumatic capillary refill sensor connected to the base unit described in Chapter 9, and a second capillary refill device using a CMOS camera with white light illumination (found to produce poor quality data and not discussed further here).

The two capillary refill devices used in the PaedCEW study were not designed to produce quantitative NIRS data. The pneumatic sensor used LED light sources at 535nm, 640nm, and 950nm wavelengths, with a source to detector separation that varied from 3 to 5mm between the photodiode and different wavelength LEDs. The camera system used diffuse white light illumination with an off the shelf RGB camera unit. The goal of the devices having been not to provide quantitative molarity data, as is the case with

NIRS, but merely to replicate the clinical capillary refill test as closely as possible, in order to substitute a measurement of refill time. In the PaedCEW study this time was defined as the time for light intensity to fall to a threshold value following the release of pressure.

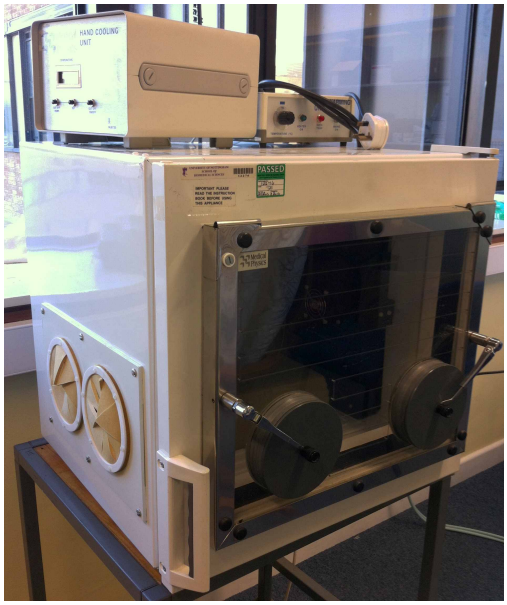


Figure 10.2: The customised refridgerator used as a chamber for cooling studies.

The extent to which the PaedCEW refill data can be used to gain a better understanding of the nature of the blood return in capillary refill is clearly limited in the absence of either any quantitative molarity output from the device or any independent and simultaneous gold standard measurement of perfusion (for example venous occlusion). The increase in capillary refill time as temperature decreases is well understood, and has been demonstrated in numerous studies, for example Gorelick et al [204], who found a significant difference in manually assessed capillary refill time between subjects measured after a 15 minute acclimatisation period in a room held at a temperature of approximately 26°C or 19°C. Watson and Kelly [238] found that ambient temperature was a significant predictor of capillary refill time in a large group of 331 healthy adult volunteers. Schriger et al [203] studied the effect of immersion in a 14°C water bath on manually assessed capillary refill time in a large group of 304 volunteers, finding a statistically significant ($p<0.01$) increase in capillary refill time (median time 1.3 to 2.9 seconds) after the water bath exposure.

Thus it would appear that the PaedCEW capillary refill data taken in isolation is of limited scientific relevance, as the size of the group, (20 adult, 17 child subjects) is much smaller than many of the existing capillary refill studies. The CRT data provide only limited new insights into capillary refill; that it can be automated successfully as part of a trial, and measured in the near infrared.

Laser Doppler and pulse oximetry sensors were used in the study, with the data recorded simultaneously with the capillary refills. This aspect of the study would appear to be novel; a literature review reveals few studies involving simultaneous measurement of capillary refill time with either laser Doppler flux or pulse oximetry derived PI, although

a 2002 study by Lima et al found capillary refill time and PI to be approximately equal in sensitivity and specificity in the prediction of high core-peripheral temperature difference [14]. Thus it can be seen that the PaedCEW data may improve understanding of the relationship between different perfusion metrics as perfusion changes, and also aid understanding of which tissue properties the capillary refill test is sensitive to. Although it would be expected that the capillary refill time should increase, and the pulse oximeter PI metric and laser Doppler flux metric fall as the forearm is cooled, with a significant population size taking part in the study, it may be possible to correlate the changes in these three different metrics and potentially better understand the changes in blood flow at the capillary level induced by the environmental temperature change, and how they relate to the suitability of the individual techniques for perfusion monitoring.

For this reason it was decided to conduct an independent analysis of the data from the PaedCEW trial, with the datasets from each volunteer being divided into three time intervals; baseline before the forearm was placed into the chamber, a cooldown period after the forearm was placed into the chamber (nominally of 15 minutes duration), and finally a second cooled "baseline" for the final 5 minutes for chamber exposure (although it was found that forearm temperature often decreased further during this period).

The Masimo Radical 7 perfusion index output was available for all the volunteers, sampled at a rate of 1Hz, and was judged to be of sufficient quality for either the mean or median value to be taken over baseline windowing intervals before and after cooling. However, the Moor VMS device data, sampled at 40Hz, were found to suffer from very significant upwards spiking, apparently caused by sensor motion. For this reason a filtering algorithm was developed for processing of the laser Doppler data. Analysis of the PaedCEW trial data follows later in this chapter (section 10.5).

10.2 Venous occlusion and NIRS

Methods for the processing of raw CW-NIRS data were discussed at length in Chapter 8. However, these were applied previously only to capillary refill data. Venous occlusion used in conjunction with NIRS was previously discussed in Chapter 6, and has been used in a large number of perfusion studies, particularly of muscle perfusion, for example by Van Beekvelt et al [164]. Typically raw near infrared intensity data have been processed via spectroscopic techniques to derive tissue haemoglobin molarity, and a linear fit made to the measured HbO_2 and Hb molarities as a function of time during the occlusion.

Attempts at numerous methods of processing continuous wave NIRS intensity data to yield tissue HbO_2 and Hb molarities in Chapter 8 met with varying degrees of success.

The differential path length method yielding perhaps the most promising results. For this reason it was decided to use the differential path length method to convert the near infrared intensity change data measured from the custom made NIRS hardware into changes in tissue molarities. Like capillary refill, venous occlusion is a differential technique, so there is no requirement for knowledge of baseline tissue Hb or HbO_2 molarity. Baseline properties are part of the differential path model itself, so it can be seen that there is an intrinsic assumption that only small changes in tissue optical properties (μ_a and μ_s') will be observed.

In an attempt to verify whether or not useful conclusions regarding perfusion could be reached using a blood pressure cuff inflated to 50mmHg applied in conjunction with the custom made CW-NIRS sensor placed on the Thenar Eminence, a series of simple tests were made using a healthy adult volunteer and a 5°C water bath.

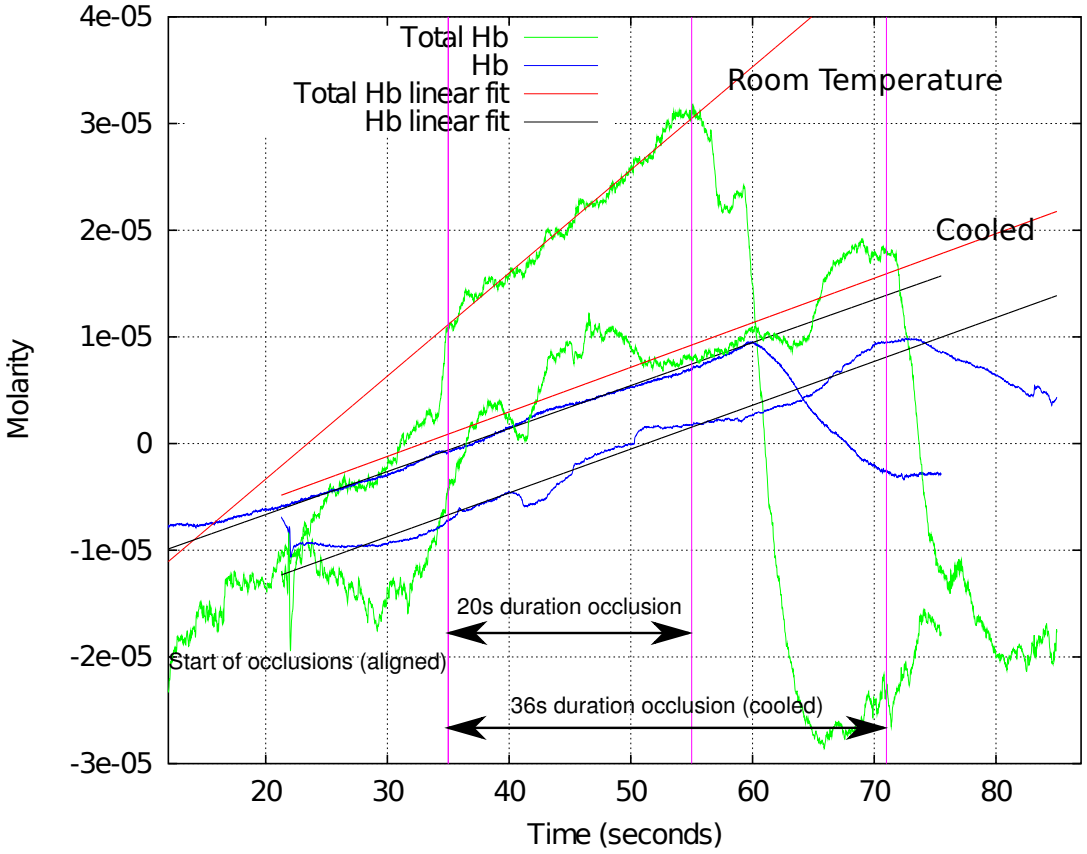


Figure 10.3: Venous occlusion results before and after water bath immersion.

Figure 10.3 shows venous occlusions of 20 seconds (before cooling) and 36 seconds duration, recorded from the palm. The NIRS sensor was placed at the same location on the Thenar Eminence for both recordings. The UCL differential path length dataset was applied to convert measured intensity changes into molarities, using the same processing steps as were applied earlier to capillary refill data, and straight line fits were

made through the region of constant 50mmHg air pressure application using a least mean squares fitting algorithm. The times at which stable pressure was reached, and the points at which pressure was released back to zero are indicated with the vertical pink lines on the plot. It should be noted that the Omron M2 unit used to inflate the cuff takes approximately 10 seconds to develop pressure, potentially explaining some of the artefacts that are observed in the signal prior to the pressure hold.

It can be seen from the results that there is no significant difference in the rate of change of Hb molarity between the two recordings. This would be the expected result, as arterial blood is around 98% saturated in healthy adults, so the inflowing blood contains a relatively insignificant amount of Hb. It can thus be seen that at low perfusion states, the Hb accumulating within the tissue is mostly a result of metabolic oxygen consumption, and as there was no muscle activity within the hand (it was stationary throughout the tests), the rate of Hb increase would not be expected to change between the room temperature (normal perfusion) and cooled (low perfusion) states. The StO_2 results from NIRS measurements of the Thenar Eminence published by Lima et al [209] would appear to support this conclusion.

There is a clear increase in the oxygen extraction index (the Hb gradient as a fraction of the total Hb gradient) between the two recordings, with a change from approximately 45% to over 90% seen from the gradients of the straight line fits. There is some agreement between the results and published studies, for example Abramson et al found forearm oxygen extraction index increased from an average of below 25% with the arm immersed in a 45°C water bath to approximately 50% in a 17 °C bath [239]. Similar results were obtained for heating only by Gallen et al [240]. However, published studies of oxygen extraction index in the hand when immersed in a low temperature water bath at around 5°C could not be found, so it is unclear if the 17°C and 45°C results can be extrapolated, but if they were to be, it would seem that the apparent close to total desaturation seen from the NIRS data is perhaps slightly excessive, and the true extraction index might be expected to be closer to 75% than the over 90% seen.

The experiment was repeated on a subsequent day without cooling the subjects forearm to test the repeatability (figure 10.4). The straight line fit to the repeated experimental data is indicated as the dotted lines. It can be seen that the rate of Hb increase was approximately 25% lower on the repeat run of the experiment. From the data prior to and after the venous occlusion it can be seen that there are significant artefacts, which may be a cause of this change, as it would seem unlikely that the basal metabolic rate of the tissue would change by as much as 25%, although some change would seem feasible, so this may be a combination of multiple biassing effects.

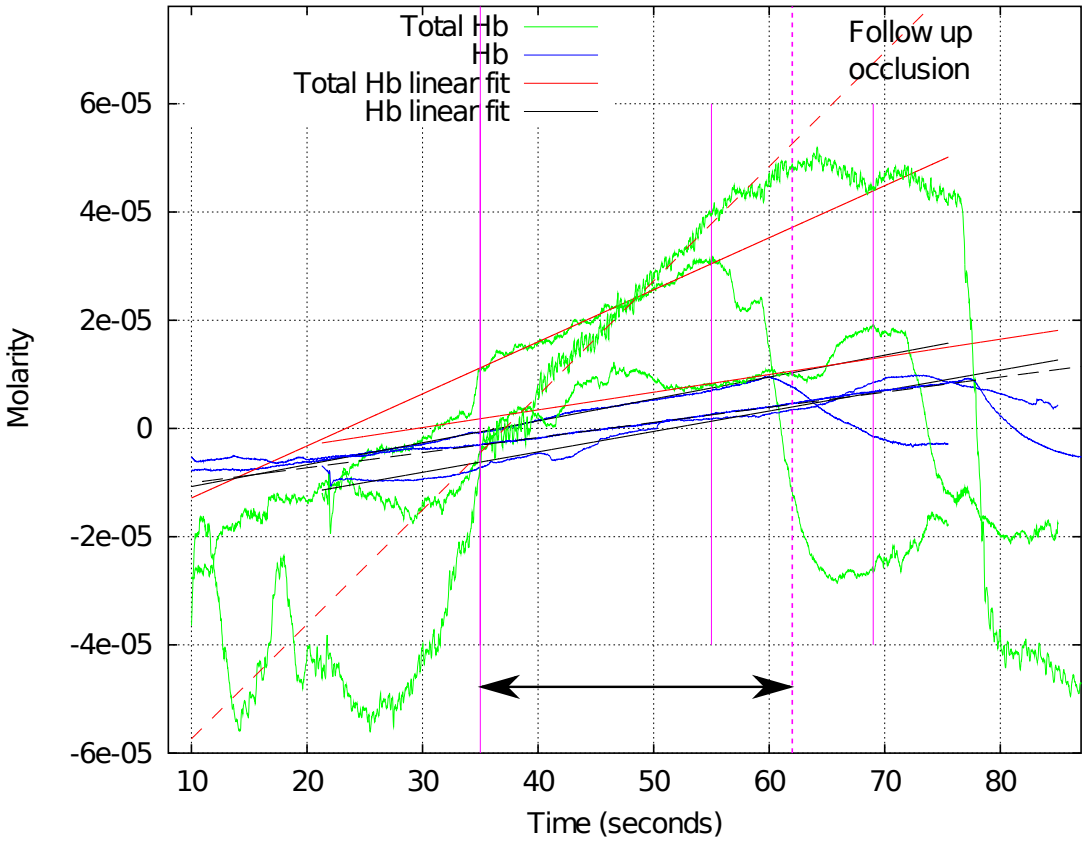


Figure 10.4: Venous occlusion results from figure 10.3 with data from a subsequent day overlaid (indicated using arrow).

The resting state perfusion from the repeated experiment (i.e. total Hb gradient) is approximately double its prior resting state value. The room temperature was noted as 18°C during the first run of the experiment, compared to 23°C during the second, so there may well have been a real perfusion increase between the two experiment runs.

With an estimate of blood haematocrit, it is possible to convert total Hb gradient during venous occlusion to a volumetric perfusion, typically expressed in $ml/100g/minute$ units. From figures 10.3 and 10.4, the rate of total Hb increase is of order $3.3 \times 10^{-7} molL^{-1}s^{-1}$ to $2.0 \times 10^{-6} molL^{-1}s^{-1}$. If the blood haematocrit is assumed to be equal to 2.35×10^{-3} (typical value for a healthy adult male), and tissue density to be $1.05 KgL^{-1}$, then the volumetric perfusion can be found:

$$P = \frac{\dot{Hb}_{total} \times 6 \times 10^3}{Hb_{blood}\rho_{tissue}} \; (ml/100g/minute). \tag{10.1}$$

Where Hb_{blood} is the total blood haemoglobin molarity, and the factor of 6×10^3 scales from SI units to a $100g/minute$ scale. Solving for the gradients seen in the venous occlusion tests gives a perfusion between 0.8 and $4.8ml/100g/minute$.

These values range from similar to several times greater than those obtained in the

calf of a group of healthy neonates by Zaramella et al [180]. However, Harrison et al found volumetric perfusion of around an order of magnitude higher than this, with an average perfusion of $52.3 \pm 16.2 ml/100g/minute$ in the forearms of healthy adult volunteers [7]. Woinarski et al found a mean forearm blood flow of $3.5 ml/100g/minute$ (with standard deviation of 1.5) in a group of 9 healthy adult volunteers using strain gauge plethysmography [167]. Woinarski also used NIRS simultaneously on the Thenar Eminence. The NIRS being measured with a proprietary device which output total haemoglobin index in arbitrary units, meaning the NIRS data could not be directly converted to volumetric perfusion. However, a strong correlation (Pearson p value $<$ 0.01) was found between NIRS perfusion measured in arbitrary units and strain gauge volumetric perfusion.

Although the volumetric perfusion values obtained from the data in figure 10.4 are within the same range as those published by Woinarski and by Zaramella, they have more variation between successive occlusions than in either of these cited publications. The precise location and depth sensitivity of the NIRS sensor may well have had a significant influence on the perfusion values which were obtained, as the source to detector separation (5mm) was significantly shorter than the 30 to 50mm separations which are typically used in reflection NIRS studies.

In the case of Harrison et al, the discrepancy is larger and thus more difficult to explain. However, it is notable that the H_2 perfusion measurements were conducted using a heated sensor at over 40°C, which would have elevated the perfusion considerably, a fact that was noted by the authors. The IAP (4-iodoantipyrine) tracer measurements from Harrison et al gave significantly lower perfusion values than the H_2, a discrepancy that suggests this conclusion may be a valid one. There are also other possibilities, for example an unaccounted for error in the IAP measurements, or perhaps that the IAP measurements were not measuring true perfusion. IAP clearance results published by Dentry et al [241] for the adult forearm would appear to be in good agreement with both the Woinarski and Zaramella studies, so it would appear that the Harrison results are anomalously high. One tentative explanation for this could be that the Harrison study was most sensitive to the perfusion in the skin surface rather than to the deeper tissue probed by NIRS techniques, or that of the entire arm probed by use of strain gauge.

In conclusion: comparing the perfusion values found from application of venous occlusion and NIRS to published studies, it can be seen that the venous occlusion results are within the expected range, apparently demonstrating that the custom made NIRS hardware is capable of reasonably accurate quantitative results. Capillary refill results were presented in Chapter 8 using the same differential path inversion technique and

hardware, yielding perfusion values that are around two orders of magnitude higher than expected. So, the relative accuracy of the NIRS hardware used with venous occlusion supports a conclusion that capillary refill is not a direct measurement of perfusion.

10.2.1 Error analysis of venous occlusion data

To further develop the venous occlusion technique, an attempt was made to estimate the non systematic error on the NIRS derived perfusion values.

Polynomial fitting was used to fit linear (i.e. first order polynomial) Hb and HbO_2 trends to the molarity data in figures 10.3 and 10.4. Error estimates can be generated for polynomial fits on the assumption that the errors caused by signal noise are random and that their probability distribution is known. However, the probability density function of the noise seen in the figures is clearly difficult to estimate, as both the underlying signal and the probability density function of the noise may be fluctuating with time (heteroscedasticity), making separation of the true signal from the noise very challenging.

Furthermore, the noise does not appear to have a uniform spectral density; from the signal fluctuations it can be seen that it is clearly not white noise, and would appear to be closer to $1/f$ noise due to the dominance of long period cycles. Numerical derivation of both the probability density function and the power spectrum is difficult in the absence of independent knowledge of the true tissue molarity versus time. Ideally an error model incorporating serially correlated disturbances would be of use here.

There are two possible approaches for estimating the autocorrelation of the noise disturbance: a relatively long baseline measurement could be made, and its autocorrelation function used as a correlated noise model. This technique involves an implicit assumption that there is no baseline shift in the tissue Hb and HbO_2 molarities during the course of the measurement period, although simple first order polynomial de-trending may be of use here. Alternatively, a fitting or analysis technique which includes a noise correlation term directly as part of the optimisation parameters could be used.

The Durbin-Watson test

Many authors recommend use of a Durbin-Watson (DW) test prior to making any decision on error estimation:

$$d = \frac{\sum_{t=2}^{t=T}(e_t - e_{t-1})^2}{\sum_{t=2}^{t=T}(e_t)^2}. \tag{10.2}$$

Where e_t is the residual from the fit to the data at time t, and the value of the Durbin-Watson coefficient d ($0 \leq d \leq 4$) in equation 10.2 indicates the amount of serial correla-

tion which is present. A value of 2 indicating no serial correlation, less than 2 indicating serial correlation and greater than 2 serial anti-correlation.

A linear least squares regression was applied to fit a linear ramp in Hb and HbO_2 molarity with time to the second venous occlusion dataset recorded from the healthy adult volunteer. The Durbin-Watson test was then applied to the attenuation residual data at each wavelength, yielding the results shown in table 10.1.

λ (nm)	d
660	1.8665e-03
750	7.4503e-04
810	1.0457e-03
870	7.4184e-04
950	8.3712e-04

Table 10.1: DW parameter d versus wavelength

The mean d value from table 10.1 is 1.04×10^{-3} with standard deviation of 4.74×10^{-4} across the multiple wavelenghts. It can thus be seen that the residual drift is highly correlated with time, making an uncorrelated noise model such as white noise inappropriate, as suggested by visual observation of the data in figures 10.4 and 10.3. It would thus seem likely that $1/f$ noise or a random walk would be a more accurate noise model to use, but the precise details of the nature of the drift are still unclear from this result.

De-trended Fluctuation Analysis

De-trended Fluctuation Analysis (DFA) was introduced by Peng et al [242], and has been applied to the analysis of time series data where an underlying trend may be present, but is hidden beneath noise of unknown properties. The DFA algorithm was later extended to the case of multiple unknown trends in the time series data, so called multi-fractal DFA, or MDFA. Unfortunately both the original DFA algorithm and the later MDFA methods only serve to indicate if an underlying trend may be present, and to determine the nature of the underlying noise (e.g. if $1/f$, Brownian motion, random walk or white noise is the best model for the noise). The precise statistical distribution of the underlying noise cannot be determined using DFA alone.

In the case of the spectrometer data, application of DFA is complicated due to the nature of the fitting problem; the raw attenuation data are recorded from five wavelength channels on the NIRS device, giving five degrees of freedom, whereas the fit is to a linear change in Hb and HbO_2 tissue molarity with time, i.e. two degrees of freedom.

It is possible to extend DFA to analysis of time series state vectors in multiple dimensions, for example Gu et al [243] describe application of a multi-fractal, multi-dimensional DFA method to recorded traffic datasets. However, in the case of the five wavelength spectrometry data the number of degrees of freedom is higher in the data (5) than the two degrees of freedom for molarity gradients of the Hb and HbO_2 with time. No description of a DFA algorithm that was directly applicable to the spectrometer data could be found in published work, but it may still be possible to apply DFA.

The original DFA algorithm classifies signals according to RMS error in polynomial fits to short intervals of time series data, taken from the full dataset. The mean RMS error across all of the intervals of the dataset is recorded as the interval width is changed, and the mean RMS fit error plotted as a function of the interval width on a log scale graph. The gradient of a linear fit to the plot is found (this is usually denoted α), and the value of this coefficient used as a diagnostic indicator as to the type of noise that is present. Values of α of around 1 are typically seen for $1/f$ noise processes, compared to 1.5 for random walk or Brownian noise processes.

It would clearly be possible to modify this polynomial fitting process to a second order fit of Hb and HbO_2 molarity ramps with time, with the RMS error now being between the calculated 5 wavelength attenuation change from baseline and the recorded spectrometry data. However, such an approach would ignore the fact that the total Hb tissue molarity will theoretically be increasing linearly with time throughout the duration of the occlusion.

An alternative approach would be to use a linear regression to fit linear trends in Hb and HbO_2 tissue molarity to the attenuation data, followed by application of DFA to the attenuation residual for each wavelength channel. The scaling exponent α would then allow the nature of the fit error noise to be determined.

Figure 10.5 shows the result of DFA analysis of the fit residuals from the second venous occlusion test carried out on the healthy adult volunteer. It can be seen that the lines are relatively straight, with gradients of around 1.5, close to the theoretical value for a random walk.

The mean α in table 10.2 is 1.5141 and standard deviation is 0.1042. Thus it is clear that the noise is very well approximated by a random walk, which would have a theoretical α value of 1.5, much less than one standard deviation away from the mean value of 1.51 (3s.f.) observed. This is an important finding, as a noise model can now be applied to the spectroscopy data, providing a framework for error analysis.

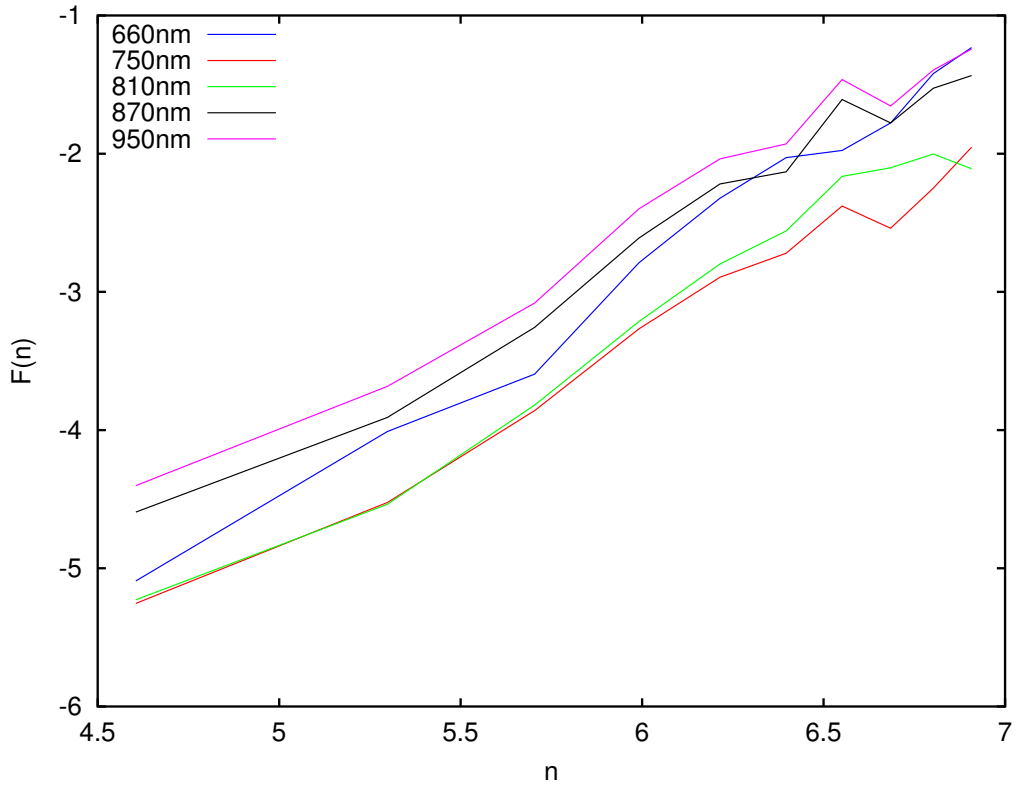


Figure 10.5: DFA analysis of the residuals at each wavelength from a model fit using linear Hb and HbO_2 ramps.

Hypotheses concerning the nature of the noisy drifts in fit error

Durbin-Watson is clearly easier to apply than DFA, but it does not offer the same degree of insight into the nature of the signal. Durbin-Watson is only applicable to serial correlation, and is not able to offer any indication of heteroscedasticity. However, considering the physical nature of the likely error terms, it is unclear if heteroscedasticity will be significant, especially when compared to serially correlated errors. More specifically, from the design of the reflectance mode spectroscopy device, which is not adhered to the skin surface and relies on a low contact pressure and glycerol contact gel to maintain coupling, it can be seen that contact errors are likely to be a significant source of error, and that contact pressure may lead to tidal motions of blood in a similar fashion to the skin blanching process of the capillary refill. These errors will be a function of subject hand tremors and control loop wander, which are likely to be Brownian and $1/f$ processes respectively.

There may also be absorption fluctuations of biological origin. There are numerous such processes with periods in the range of tens of seconds, for example respiration induced tidal motion of blood, respiratory mean blood pressure fluctuations, and autoregulatory

λ (nm)	α
660	1.6896
750	1.4482
810	1.5290
870	1.464
950	1.4391

Table 10.2: DFA fit α versus wavelength

baro-receptor induced blood pressure oscillation. Auto-regulatory processes lead to a 0.1Hz oscillation in mean arterial pressure and tissue blood volume fraction health adults (often termed "Mayer waves"), a phenomenon that is not fully understood currently, with some debate concerning the origin and relative importance of the various feedback paths. Malpas [244] has published an overview of current understanding of baro-receptor oscillation in both humans and animal models.

With frequency peaks in the 0.1Hz range, these oscillatory signals will clearly appear highly correlated at the spectrometer's sample rate of over 100Hz. The typical duration of a venous occlusion is less than a minute, as ideally pressure should be released from the occluding cuff before accumulating blood causes a sufficient increase in capillary bed pressure that the arterial inflow decreases. If this is allowed to occur it causes the rate of increase of Hb_{total} molarity to fall, rendering the results inaccurate. Unfortunately this allows only a few 0.1Hz oscillations to fit into any single venous occlusion, and the resultant aliasing effects may place a significant limit on the perfusion accuracy and repeatability that can be achieved when using venous occlusion.

The limited duration of the venous occlusion test carried out on the volunteer also places limits on the applicability of the DW and DFA tests, but the fact that the DFA α value was so close to 1.5 suggests that a random walk model is an accurate model over the duration of the occlusion, despite the fact that other, long period oscillatory effects may in fact be responsible for a significant proportion of the fit error.

Feasible Generalized Least Squares

Feasible Generalized Least Squares (FGLS) is a development of Generalized Least Squares (GLS) where the measurement covariance matrix is unknown. Iterative techniques, such as that described by White et al [245], are used to produce an estimated covariance matrix based on multiple GLS residuals. However, although GLS allows an arbitrary measurement covariance matrix to be used, the majority of the FGLS iteration procedures in use assume a diagonal covariance matrix, i.e. there is assumed to be heteroscedas-

ticity but no serially correlated disturbances, and thus the GLS model is simplified to a weighted least squares fit (WLS).

Two FGLS estimators which do allow serially correlated error terms, and that have seen widespread adoption are the Cochrane-Orcutt [246], and Prais-Weston [247] estimators. These estimators rely on noise autocorrelation models with reduced degrees of freedom in order to make a solution feasible.

Newey-West standard errors

In the presence of serially correlated noise and heteroscedasticity, OLS will not be the so called Best Unbiassed Linear Estimator (it is not "BLUE"). But, it is still unbiassed, meaning that the expectation value of the fit parameters will not change under these conditions (i.e. it is Unbiassed), although the standard deviation of the fit parameters across multiple experiments may not take the lowest possible value (i.e. it is not the Best Estimator).

Thus, even in conditions where OLS is not the best estimator, it is possible to apply it providing it is still unbiassed. At the expense of a decrease in performance, it can then be followed by a separate procedure for error modelling. However, this will result in an increase in the computed standard errors above those of the "best" estimator (in practice this effect is usually found to be small).

The Newey-West estimator [248], a development of the earlier Cochrane-Orcutt estimator [246], is the most widely used such procedure. Newey-West assumes prior application of an ordinary least squares fitting algorithm, with the residuals from this fit being taken as input parameters. Under the assumptions of a power law serial correlation function and/or heteroscedastic errors, the Newey-West algorithm can accurately determine the covariance matrix of the fit parameters from the OLS residuals alone.

Application of Newey-West to spectroscopy data is not possible without modification of the multiple wavelength attenuation data, as the Newey-West algorithm requires a vector argument of OLS residuals. A preprocessing technique to reduce the dimensionality of the OLS residual data is therefore required before the fit errors can be estimated. One simple approach would be to stack the attenuation data for the multiple wavelengths, with each one following the previous along a single vector. Unfortunately this simple approach would be likely to lead to significant biassing of the autocorrelation function, as there would be significant transients in the residual data at the transition point between each wavelength. However, if this same technique is used in conjunction with zero padding, it is possible to mitigate the effect of the transients. By adding a zero padded

region of length L (the Newey-West maximum correlation length), the possibility of transient induced biassing of the autocorrelation is eliminated at the expense of a slight downward biassing of the error autocorrelation estimation. This preprocessing algorithm is illustrated schematically in figure 10.6

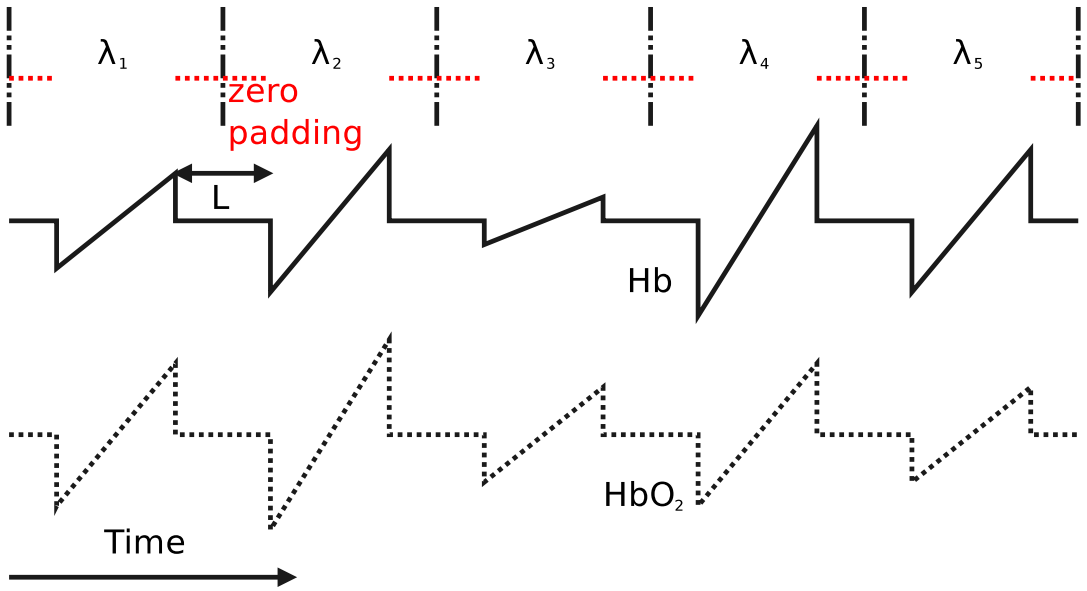


Figure 10.6: Outputs from the preprocessing algorithm used prior to Newey-West, L is the lag limit (see text for details).

The Hb and HbO_2 ramp functions are passed to an OLS regression algorithm together with the zero padded attenuation data (these are illustrated as the solid line, the dotted line and the λ_1 to λ_5 regions in figure 10.6). So as to avoid any requirement to preprocess attenuation data to ensure biassing does not occur due to differing mean attenuation between the wavelength channels, the Hb and HbO_2 molarity ramps are offset to give a mean of zero within each wavelength range. This allows a simple data processing system to be used for the generation of attenuation change from raw intensity, for example simply taking the mean intensity value for each channel followed by application of equation 10.3

$$A_n = -\ln \frac{I_n}{< I >}. \tag{10.3}$$

Where $< I >$ is the mean intensity for a given wavelength channel.

It should be noted that although this approach does not completely ignore the possibility of baseline drift (which may arise due to variable optical coupling with the skin surface, or LED output power fluctuations with temperature for example), the AR(1) noise model as used in Newey-West is unlikely to be a good model of such drift, particularly in cases such as LED output drift, where a linear model is more likely to be accurate. For this

reason, use of the periods of baseline before and after the venous occlusion to fit a linear trend to the intensity data at each wavelength may result in superior performance to Newey-West alone.

Error propagation modelling

When used in combination with NIRS, venous occlusion is capable of providing two metrics; the venous oxygen saturation (Sat_v), and the volumetric total haemoglobin molar perfusion ($Perfusion_{Hb}$). These can be calculated from the rates of increase of Hb and HbO_2 molarity ($rate_{Hb}$ and $rate_{HbO_2}$) through application of the following equations:

$$Perfusion_{Hb} = rate_{Hb} + rate_{HbO_2}, \tag{10.4}$$

$$Sat_v = \frac{rate_{Hb}}{rate_{Hb} + rate_{HbO_2}}. \tag{10.5}$$

Application of the preprocessor followed by the Newey-West algorithm to the OLS residual data gives an estimated error covariance matrix for the two molarity rate values. Introducing errors in $rate_{Hb}$ and $rate_{HbO_2}$ as e_1 and e_2 respectively, the covariance matrix is defined as:

$$\begin{pmatrix} < e_1e_1 > & < e_1e_2 > \\ < e_2e_1 > & < e_2e_2 > \end{pmatrix}.$$

As an estimate for this matrix is provided by the Newey-West algorithm, it is possible to use small error assumptions to propagate the e_1 and e_2 error estimates through equations 10.4 and 10.5, giving covariance estimates for the perfusion error e_p and venous saturation error e_v.

The perfusion error covariance is easily calculated from the $e_{1,2}$ terms:

$$< e_pe_p >=< (e_1 + e_2)(e_1 + e_2) >=< e_1e_1 > + < e_2e_2 > +2 < e_1e_2 > . \tag{10.6}$$

The venous saturation covariance is somewhat more complex, and a Taylor expansion is required:

$$< e_ve_v >= \left(\frac{\partial e_v}{\partial e_1}\right)^2 < e_1e_1 > + \left(\frac{\partial e_v}{\partial e_2}\right)^2 < e_2e_2 > +2\frac{\partial e_v}{\partial e_1}\frac{\partial e_v}{\partial e_2} < e_1e_2 >. \tag{10.7}$$

Where the first partial derivative is:

$$\frac{\partial e_v}{\partial e_1} = \frac{\partial Sat_v}{\partial rate_{Hb}} = \frac{\partial \frac{rate_{Hb}}{rate_{Hb}+rate_{HbO_2}}}{\partial rate_{Hb}} = \frac{1}{rate_{Hb}+rate_{HbO_2}}\left(1 - \frac{rate_{Hb}}{rate_{Hb}+rate_{HbO_2}}\right), \tag{10.8}$$

$$\frac{\partial e_v}{\partial e_1} = \frac{1 - Sat_v}{Perfusion_{Hb}}, \tag{10.9}$$

and the second is:

$$\frac{\partial e_v}{\partial e_2} = \frac{\partial Sat_v}{\partial rate_{HbO_2}} = \frac{\partial \frac{rate_{Hb}}{rate_{Hb}+rate_{HbO_2}}}{\partial rate_{HbO_2}} = \frac{-rate_{Hb}}{(rate_{Hb}+rate_{HbO_2})^2}, \tag{10.10}$$

$$\frac{\partial e_v}{\partial e_2} = \frac{-Sat_v}{Perfusion_{Hb}}. \tag{10.11}$$

Putting these together a final expression is obtained:

$$< e_v e_v >= \frac{1}{Perfusion^2}\left((1 - Sat_v)^2 < e_1 e_1 > + Sat_v^2 < e_2 e_2 > + 2Sat_v(Sat_v - 1) < e_1 e_2 >\right). \tag{10.12}$$

Thus through equations 10.6 and 10.12 it is possible to calculate estimated venous saturation and haemoglobin perfusion errors from the OLS fit parameters and Newey-West estimated covariance matrix.

From the total haemoglobin perfusion and venous oxygen saturation it is possible to calculate a third tissue parameter, the metabolic rate of oxygen, provided that the arterial oxygen saturation is known. Arterial oxygen saturation can be found using a conventional pulse oximeter, or alternatively estimated as being equal to 98%, the value in healthy adults and children is within one to two percent of this. The metabolic rate of oxygen is then defined simply as the total Hb molar perfusion multiplied by the product of the difference between arterial and venous saturation and the oxygen carrying capacity of haemoglobin, which is accuratly known from published experimental data:

$$Met_{O_2} = (S_a - S_v)C_{Hb}Perfusion_{Hb}. \tag{10.13}$$

Defining e_M as the error in the metabolic rate of oxygen, it follows that:

$$< e_M e_M >= Met_{O_2}^2(< e_v e_v > + < e_p e_p > + < e_a e_a >). \tag{10.14}$$

Where e_a is defined as the error in the arterial oxygen saturation (which will have an expectation value of order 1.5% for a commercial pulse oximeter and normal oxygen saturations of around 98%).

For these covariance expectation values to be converted to standard errors, a form must be assumed for the error probability density function (the PDF). Assuming a Guassian probability distribution, i.e.

$$PDF = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}. \tag{10.15}$$

Where PDF is a probability density function centred on μ with variance σ^2. The standard error is usually defined as the 95% confidence limit, meaning that from the cumulative probability function ($\int PDF dx$), standard error is found to be 1.96σ i.e. $1.96\sqrt{< ee >}$ for an error term e.

10.2.2 Analysis of the healthy volunteer data

The ordinary least squares linear fitting using the ramp function illustrated in figure 10.6 was applied to the first set of healthy volunteer occlusion data. With the built in GNU-Octave ordinary least squares function ("ols"), and a custom implementation of the Newey-West algorithm based on the original publication being used.

Experiment	Perfusion (ml/100g/min)	Venous saturation (%)	Metabolic rate (mlO_2/100g/min)
Baseline test 1	3.42 ±0.29	69.8 ±3.6	0.242 ±0.015
Cooled test	1.02 ±0.20	-8.4 ±24.5	0.259 ±0.023
Baseline test 2 (4 days later)	1.40 ±0.29	23.7 ±19.7	0.250 ±0.027

Table 10.3: Fit results and errors from the first healthy volunteer occlusion.

The fact that the metabolic rate of oxygen uptake does not change within the margin of error is encouraging; it would be expected to remain approximately constant, as the forearm was at rest for all the occlusion tests. In a study of skeletal muscle metabolism, Zurlo et al used mechanical plethysmography combined with invasive blood sampling to measure the metabolic rate of oxygen consumption in the forearms of a group of healthy adult volunteers [249], and found similar values to those in table 10.3, with the consumption clustering around 0.14 mlO_2/100g/min and the largest value being approximately 0.17 ± 10%, thus putting the upper bound 25% below the lowest value recorded using the NIRS system. To assess the significance of this result, it may well

be relevant that the Zurlo experiment used volume plethysmography over the entire forearm and invasive sampling of a deep antecubital vein, whereas the NIRS device was sampling only a small volume of the palm.

It is also important to note that the error estimates produced by the Newey-West algorithm are not measurements of systematic error, but rather of random error estimated from the deviation between the model and the attenuation data. There are likely to be significant errors in the estimated path lengths, errors which will not be modelled by the Newey-West plug in procedure. Considering that the pathlength factor is found to vary between individuals by $\pm 20\%$ or more (for example as discussed in Kohl [148]) it should thus be unsurprising that the metabolic rate as derived via NIRS is somewhat higher than would expected from more accurate invasive studies.

For comparative studies of perfusion changes, this systematic error may not be a serious problem, as if the measurement site remains unchanged, the path length through the tissue would not be expected to change significantly.

10.3 Capillary refill and NIRS

The proposed experimental protocol illustrated schematically in figure 10.1 involves multiple capillary refills conducted before and after the forearm is cooled, and repeated on each occasion both with and without simultaneous application of arterial occlusion. Earlier, proof of concept, capillary refill experiments in Chapter 8 demonstrated that a complete capillary refill observation cycle (press, release, NIRS observation of the return to baseline), can be completed in less than 10 seconds. This means that a total of three refill cycles could be completed both before and after arterial occlusion with the 30 second intervals shown in figure 10.1.

Unlike venous occlusion, where the theoretical explanation of local blood pooling is well accepted, and was found to provide a good match to the experimental data analysed earlier in this chapter, the precise source of the blood returning to fill the capillary bed is unclear. This complicates error analysis, as for a technique such as Newey-West to be applied, the raw attenuation data need to have been fitted to a model and a set of fit residuals produced.

Alternatively, as there are multiple capillary refill measurements made during each 30 second observation window, it is feasible to use the standard deviation to produce an error estimate. Theoretically, assuming a Guassian PDF for the capillary refill metric (which may be taken as either the time to return to baseline or the total Hb molarity gradient following pressure release), then the square root of the variance (σ) is equal to

the corrected sample standard deviation (s) by definition for a large number of samples. The corrected sample standard deviation is defined as follows:

$$s^2 = \sum_{n=1}^{N} \left[\frac{(x_n - \bar{x})^2}{n-1} \right]. \tag{10.16}$$

If the root of the sample variance σ is known, under the assumption that the samples are not serially correlated, the variance of the mean can be easily calculated as:

$$Var(\bar{x}) = \frac{\sigma}{\sqrt{n}}. \tag{10.17}$$

For a finite number of samples, n, s is no longer equal to the root of the variance, σ, and the sample standard deviation is found to be an underestimate of the true sample variance. Cochrane's theorem [250] leads to a correction term $c_4(n)$ which can be used to scale the measured standard deviation to reach a best estimate for the sample variance based on the number of samples n:

$$< s >= c_4(n)\sigma, \tag{10.18}$$

where $c_4(n)$ is found to have the following values:

n	$c_4(n)$
2	$\sqrt{\frac{2}{\pi}}$
3	$\frac{\sqrt{\pi}}{2}$
4	$2\sqrt{\frac{2}{3\pi}}$
5	$\frac{3\sqrt{\pi}}{4\sqrt{2}}$
6	$\frac{8}{3}\sqrt{\frac{2}{5\pi}}$

Table 10.4: $c_4(n)$ values versus n.

It is important to note that equation 10.18 only describes the expectation value of the standard deviation of the samples. As the number of samples decreases, the uncertainty on the true value of s increases. The variance of s has been found to be described by the following expression (see Lehmann, p31 [251]):

$$\sqrt{Var(s^2)} = \sigma^2\sqrt{\frac{2}{n-1}} = \frac{< s >^2}{c_4(n)^2}\sqrt{\frac{2}{n-1}}. \tag{10.19}$$

In the case of three samples, i.e. $n = 3$, this can be simplified to

$$\sqrt{< Var(s^2) >} = s^2\frac{4}{\pi}, \tag{10.20}$$

meaning that the error on s^2, and therefore on σ^2, is larger than s^2 itself. However, this does not mean that no useful inference can be drawn from as few as three capillary refill cycles, as it is still possible to place a likely upper bound upon the refill time variance σ. The number of capillary refills in the standard deviation calculation could be increased to six if the refills conducted both before and during arterial occlusion were used. However, it is not clear if the variance of the capillary refill time is fixed, or if it scales with the refill time itself, which would make such a calculation an invalid estimator for s.

If the number of samples was to be increased to 6, this would lead to a lower uncertainty in the standard deviation. From application of table 10.4:

$$\sqrt{< Var(s^2) >} = s^2 \frac{45\pi}{128} \sqrt{\frac{2}{5}} = 0.70 s^2 (2s.f.). \tag{10.21}$$

This is almost half the variance on s for 3 samples, but is still sufficiently large than an accurate assessment of σ is still impossible from six samples.

The experimental protocol could be adjusted based on this finding, to increase the number of capillary refill cycles both before and after the water bath exposure, and thus obtain a more accurate estimate of the variance. However, this would increase the risk of biassing of the results due to the arm temperature increasing following withdrawal from the water bath, meaning that the venous occlusion test would be more likely to take place in a differing vascular state to the capillary refills. There is also a risk of injury to the volunteer's arm due to prolonged loss of blood supply if a large number of refills were to be conducted under arterial occlusion. For these reasons it was previously considered that the number of capillary refills be limited to six following the water bath (i.e. three with no cuff applied followed by a further three with the arterial occlusion cuff inflated above the elbow). However, it was judged feasible to increase the number of refills conducted before placing the forearm in the water bath to ten prior to the application of the arterial cuff, followed by six with the cuff applied and inflated. With ten refills, the error on the standard deviation is decreased further, to $\sqrt{< Var(s^2) >} = 0.50 s^2 (2s.f.)$.

10.4 Masimo pulse oximeter perfusion index

The so called "Perfusion Index", or PI, is a metric output from commercial pulse oximeters from Masimo Inc. (and others). This was discussed previously in Chapter 6, and operates based on the principle that a greater perfusion will tend to correlate with increased pulsatile modulation depth from the pulse oximetry sensor. To provide a third perfusion metric, it was decided to record the PI using a Masimo "Radical 7" pulse

oximeter clipped onto the fingers throughout the duration of the capillary refill experiments. Unfortunately the Masimo pulse oximetry sensors are not waterproof, so for the pulse oximetry data to be continuous during water bath cooling, pushing the entire arm into a long, thin, plastic bag was found to be a practical means of preventing ingress of water.

10.4.1 PI baseline and breath hold experiments

The Radical 7 device has an RS232 serial output streaming heartrate, oxygen saturation, and perfusion index data at 1Hz. This was used to log baseline data from two healthy adult volunteers in an attempt to understand the characteristics of both the filtering carried out within the device itself (and thus the transient behaviour after the pulse oximeter is first turned on), and also of any noise or oscillatory components (both of random nature and any of biological nature). This baseline experiment was conducted twice for each volunteer.

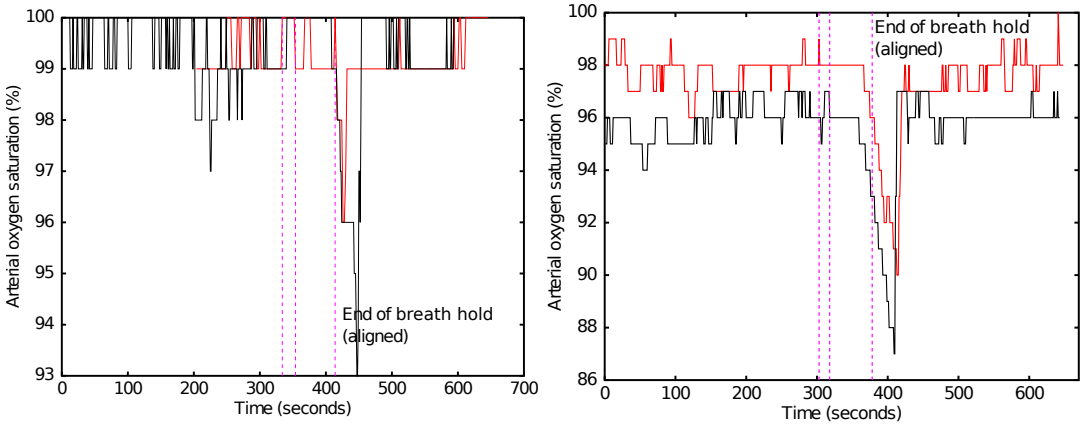


Figure 10.7: Two pulse oximeter recordings from each of the two volunteers, showing blood oxygen saturation. Black trace is the 80 second hold, red the 60 second. Volunteer one at left, two at right.

In order to investigate the effect of breath hold on perfusion index, heart rate, and arterial oxygen saturation, the volunteers were asked to hold their breath for one minute during the first baseline recording, and 80 seconds during the second. The two recordings have been aligned to the end of the breath hold on figure 10.7, illustrated by the last vertical dashed line. The other two lines between 300 and 400 seconds are the starting times of the 80 second and the 60 second breath hold.

It can be seen that the arterial blood oxygen saturation drops significantly following the resumption of breathing, with the 80 second hold resulting in the largest decrease in saturation. An initially unexpected observation in the first volunteer was that oxygen saturation did not drop until they resumed breathing. This does not appear to be an

artefact of the response time of the pulse oximeter, as the timing of the saturation drop is identical between the two recordings taken from the first volunteer. One possible explanation is that the first volunteer was only able to voluntarily hold their breath up to the point when their blood oxygen saturation began to drop precipitously. However, this is also highly unlikely to be the explanation, as the volunteer was requested before the experiment to hold their breath for 60 and 80 seconds for the first and second recordings respectively, and they reported that they were normally able to hold their breath for considerably longer than 80 seconds. Volunteer 2 was only able to hold their breath for 75 seconds, and their oxygen saturation can be seen to drop before they started breathing, so this explanation would appear to explain the oxygen saturation behaviour seen in the second volunteer.

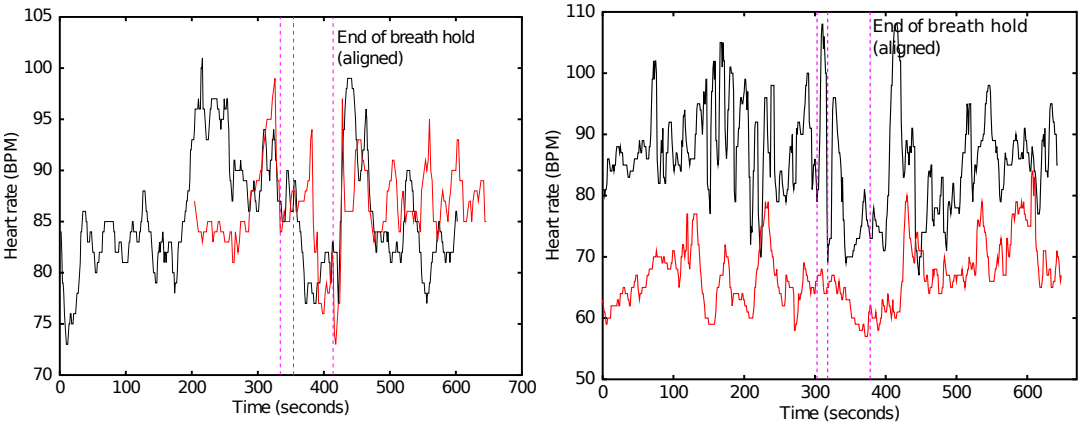


Figure 10.8: Two pulse oximeter recordings from each of the two volunteers, showing heart rate in BPM. Black trace is the 80 second hold, red the 60 second.

Figure 10.8 shows the heart rate data from the same two volunteers. The heartrate is quite variable in all four runs of the experiment, making definitive conclusions impossible. It may be notable that in two cases heartrate spiked upwards prior to the subject holding their breath, suggesting that nervousness may have contributed to the variability. It may also be of note that there is a downward trend in heartrate throughout the breath hold in all four cases. One potential explanation for this is that bio-regulatory feedback processes act to restrict perfusion under conditions of apnoea, resulting an a decrease in heart rate due to lower cardiac output. Such a decrease in perfusion might also at least partially explain the observed high arterial oxygen saturation throughout the breath hold, as decreased peripheral perfusion could help to maintain arterial saturation despite oxygen availability decreasing.

The explanation that changes in peripheral perfusion are occurring during the breath hold is supported by the PI metric data, shown in figure 10.9. In both traces there is a general trend downwards as the subjects hold their breath, followed by a recovery and

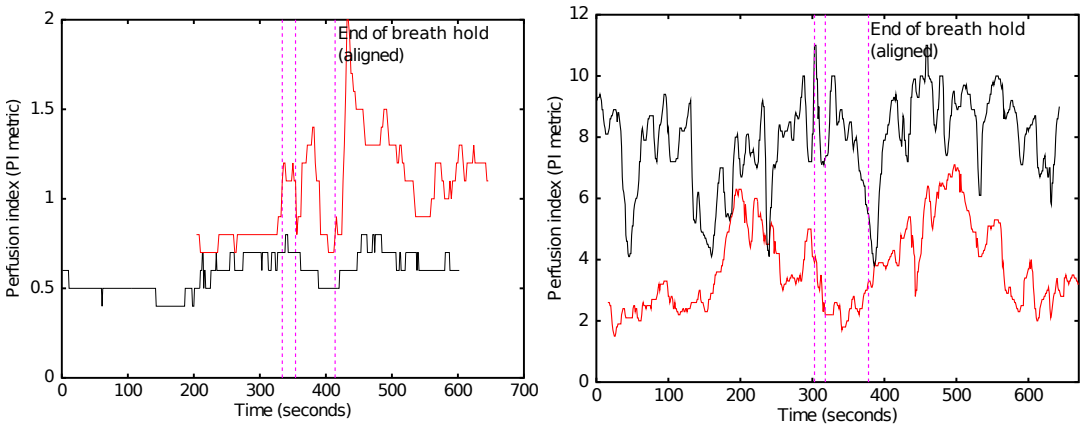


Figure 10.9: Two pulse oximeter recordings each of the two volunteers, showing the perfusion index metric output from the Masimo Radical 7 pulse oximeter. The black trace is the 80 second hold, red the 60 second.

overshoot above baseline as they resume breathing, suggesting a hyperaemic response. However, as with the heartrate traces there is considerable drift and noise, making interpretation difficult. It may be of relevance that there is apparently lower noise evident in three of the traces for the first 150 seconds or so after the pulse oximeter is first turned on. Hypothetically this could be explained by filtering processes used to estimate the pulsatile modulation depth. For example, if a rolling average of the past 150 seconds of modulation depth is used to estimate the PI metric after the point 150 seconds into the recording, then a different and more agressive filtering technique may be used for the first 150 seconds of PI output. It should be noted that the Radical 7 device was found to require the pulse oximetry sensor to be attached to the volunteer for at least ten seconds before it began to output data over RS232, suggesting that some form of rolling average with a short averaging length followed by an aggressive low pass filter may explain the behaviour of the PI output in the first few minutes of operation. Such an averaging period is suggested by O'Donnell et al in a 2005 study [252].

Extensive studies have been conducted on the influence of sleep apnoea on the cardiovascular system, for example by Imadojemu et al [253], who in a 2001 study used a laser Doppler system to investigate the effect of sleep apnoea on peripheral perfusion. This revealed a small increase in perfusion throughout the apnoea, with a large peak following the end of each apnoea, results which are suggestive of a hyperaemic response, but not of a flow shutdown during the apnoea. In his study, there was found to be no significant change in blood flow measured using Doppler ultrasound across the brachial artery. Blood pressure was found to fall during the initial stages of each apnoea, and the arterial saturation was relatively constant throughout this period, findings which are consistent with the pulse oximetry results, although not supportive of the reduced

peripheral perfusion hypothesis, as the PI metric would be expected to increase with arterial blood pressure, whilst perfusion will be a function of blood pressure and peripheral vascular resistance.

In a 1992 study of sleep apnoea [254], Garpestad et al found a decrease in cardiac output post apnoea. The arterial oxygen saturation was found to decrease slowly towards the end of each apnoea, followed by a more rapid drop once breathing resumed. As the apnoeas presented in the publication are periodic, it is difficult to draw definitive conclusions concerning the causal relationship between the various effects. However, Garpestad did see a significant increase in heart rate post apnoea, with a decrease in stroke volume being found to lead to the overall decrease in cardiac output. Thus the cardiovascular effects seen by Garpestad et al are also approximately consistent with the pulse oximetry results.

However, it should be noted that sleep apnoea is usually of relatively short duration, and the results presented in the Imadojemu publication also show cyclic apnoeas, so it may be difficult to determine the precise nature of the apnoea induced circulatory changes from the study. It may further be significant that the volunteers were lying asleep in a supine position during the Imadojemu study, as opposed to sitting in a conscious state whilst voluntarily holding their breath. Imadojemu et al conducted voluntary conscious apnoeas as part of their study, but the results were only sufficient to conclude that blood pressure rose post apnoea. Imadojemu et al concluded that the postapnoea recovery is accompanied by transient peripheral vasoconstriction. Unfortunately this phenomenon would not appear to be of use for characterisation of non invasive perfusion sensors, as the limb blood flow was found to remain constant due to the increase in blood pressure.

Nevertheless, if the conclusions reached by Imadojemu et al are correct, the breath hold tests might serve to indicate an important flaw in the use of PI as a perfusion indicator, as the PI metric would appear to indicate a perfusion fluctuation occurring under conditions where according to the results of Imadojemu et al, perfusion is likely to remain constant. It would thus appear to be likely that PI is sensitive to mean arterial pressure (MAP) either as well as or even instead of perfusion. This would be expected from a theoretical perspective, as high MAP usually leads to a larger difference between systolic and diastolic blood pressures. The relationship

$$MAP = \frac{(2diastolic) + systolic}{3} \tag{10.22}$$

is commonly employed in clinical practice as a first order approximation. This can be rearranged to form

$$MAP = diastolic + \frac{systolic - diastolic}{3}, \tag{10.23}$$

where $systolic - diastolic$ is the pulse pressure responsible for the pulsatile absorption fluctuations measured via pulse oximetry. Usually the diastolic pressure is found to be approximately proportional to the MAP, and thus the pulse pressure also changes proportionally.

This would suggest that the pulsatile modulation depth as measured with the pulse oximetry sensor will be a function of MAP, as well as of blood vessel density within the tissue, and mean vessel wall elasticity. The perfusion, a function of MAP and peripheral vascular resistance (PVR), will thus be expected to correlate with the PI metric. It may also be that decreased vessel wall elasticity is correlated with increased perfusion. Intuitively this would be expected, as vasoconstriction leads to greater vessel wall thickness. Therefore, the breath holding results may show a flaw in the use of PI as a perfusion metric, with the PI during breath hold apparently showing a decrease in perfusion followed by a recovery, as the published research is not suggestive of such a relationship. It would further appear from the published apnoea studies that a simultaneous increase in MAP and PVR could be the source of this misleading PI fluctuation.

To investigate the feasibility of this explanation for the apparent PI fluctuation during the breath holding exercise, it was repeated on the first volunteer together with venous occlusion/NIRS. As venous occlusion can only be applied for approximately 1 minute before the rise in venous pressure starts to cause a considerable decrease in perfusion, a single one minute occlusion was conducted using a blood pressure cuff whilst the volunteer held their breath. Three further venous occlusions were also made, one before the breath hold exercise, one immediately following, and a final measurement was made three minutes after the breath hold. The CW-NIRS system developed in Chapter 8 was used, with the sensor placed on the Thenar Eminence.

Occlusion	Perfusion (ml/100g/minute)	Confidence interval (95%)
Baseline 1	1.38	±0.21
Breath hold	1.38	±0.12
Post breath hold	1.71	±0.16
Baseline 2	1.43	±0.14

Table 10.5: Perfusion values recorded using venous occlusion before, during and after voluntary breath hold.

From table 10.5, there was no significant change in perfusion whilst the participant held their breath. However, when the participant resumed breathing, the perfusion increased

significantly. Such an increase would be expected to occur due to reactive hyperaemia.

The statistically significant increase in perfusion when breathing resumes suggests that breath holding exercises may offer an alternative route for inducing perfusion changes for sensor characterisation and evaluation. However, although statistically significant, the fractional increase in perfusion seen in this volunteer was only 24%, and the second baseline measurement suggests that it is a relatively transient phenomenon of short duration. In comparison, the earlier water bath cooling experiment achieved a reduction in perfusion down to 30% of its baseline value (see table 10.3) in the same volunteer, and could be sustained for an indefinite period using cooling of the forearm. For these reasons it was decided not to pursue apnoea induced perfusion changes as a means of device characterisation.

In conclusion, the PI data obtained from the Radical 7 appear to indicate a repeatable perfusion decrease at the start of the breath hold test in both of the healthy adult volunteers, whereas published studies (with a large sample size), and also a single venous occlusion test carried out on one of the volunteers show no perfusion change. This finding may demonstrate a flaw in the use of pulse oximetry PI as a measure of perfusion, but more data are needed to draw definitive conclusions due to the unstable behaviour of the PI metric.

10.4.2 PI motion experiments

A second experimental study of the PI metric was made using a single healthy volunteer. The subject sat motionless at a desk for 12 minutes, making two series of arm movements, each lasting approximately 100 seconds. In an attempt to eliminate any filtering transients or sensor acclimatisation effects, the Radical 7 device was left running with the sensor placed on the participants finger for ten minutes before the experiment.

In figure 10.10, the two intervals of limb motion are indicated using vertical dashed magenta lines. During the first interval the subject was instructed to make small arm movements of approximately 10cm displacement and one second period. During the second interval, the participant made large arm movements of approximately 50cm displacement with the same period. Different effects can be seen; during the interval of light motion it would appear that the PI may have been biassed downwards sightly, and significant short term wobble is seem, whereas the PI value is constant throughout most of the second interval. This behaviour is indicative of some form of PI measurement error estimation and adaptive filtering being used internally on the Radical 7 device. A Kalman filter would give rise to these effects and the higher PI stability during the first few minutes of device operation. If this is the case, improvements to the device firmware

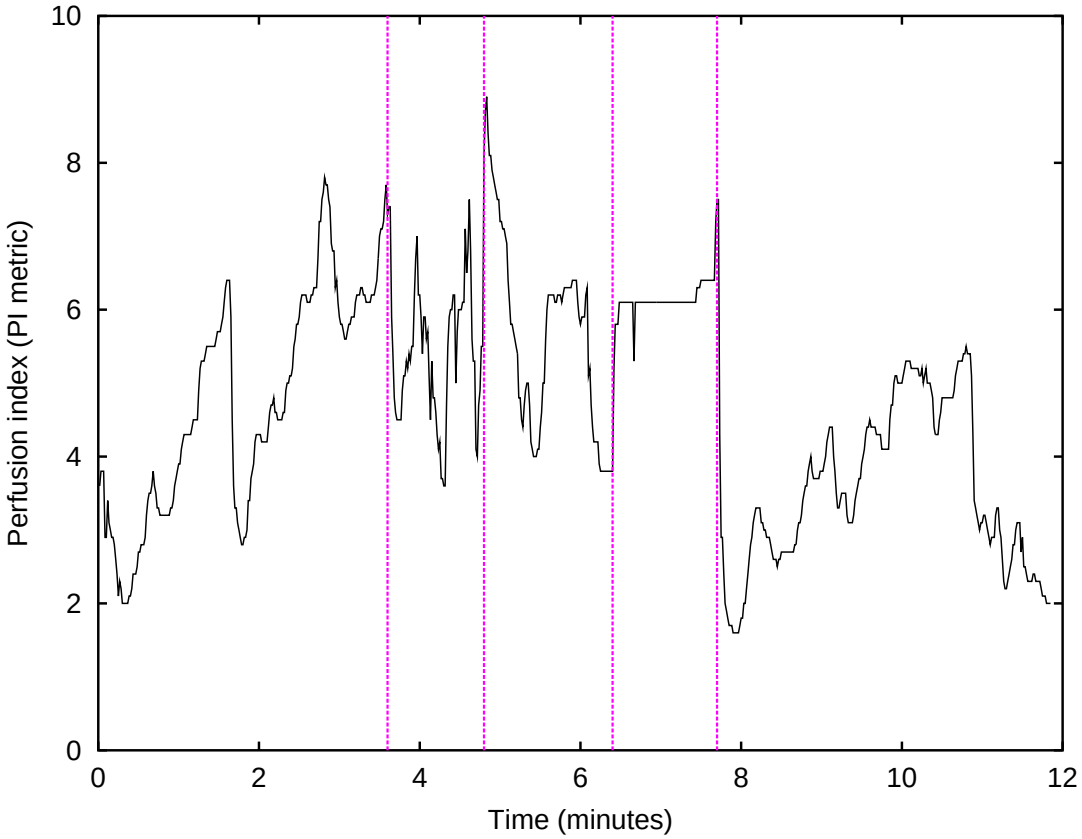


Figure 10.10: Pulse oximeter recording from healthy adult volunteer, showing influence of limb motion on PI metric.

may allow PI error estimates to be output alongside PI metric in future.

10.5 Results from the PaedCEW study

10.5.1 Moor Instruments MoorVMS laser Doppler data

The raw data saved by the MoorVMS device were found to suffer from interference in the form of spikes. These were overwhelmingly in a positive direction (i.e. rising above the baseline), and would appear to be caused by motion of the sensor relative to the skin surface, as very significant upward spiking was seen on all the PaedCEW healthy volunteer datasets for a period of around one minute as the forearm was pushed through the foam iris draft seals into the cooling chamber.

To reduce the impact of this interference, all data-points with a value outside the range of one to 60 (PU units) were replaced by linear interpolation of the nearest neighbours. A rolling window quartile (or 0.25 quantile) filter was then applied with a window length of 4 seconds, with the lower quartile value being passed to a low pass filter. Figure 10.11 shows an example of a typical recording (recording V14C) processed in this manner.

In order to show the full dynamics of the signal, a 1Hz low pass second order Bessel filter was applied following the quartile filter, meaning the oscillatory nature of the laser Doppler signal is apparent, but that longer term drift is somewhat obscured. The first 15 seconds of data have not been plotted, as the Bessel filter had a setting time of this order. Although this filtering scheme is somewhat ad-hoc, the overall design was based upon the process described in the manual for the Moor Instruments MoorVMS device, where a similar process is described as being used to produce both the value displayed on the LCD display on the front panel of the instrument (in arbitrary PU units), and the value displayed in the GUI software shipped with the device for live data viewing.

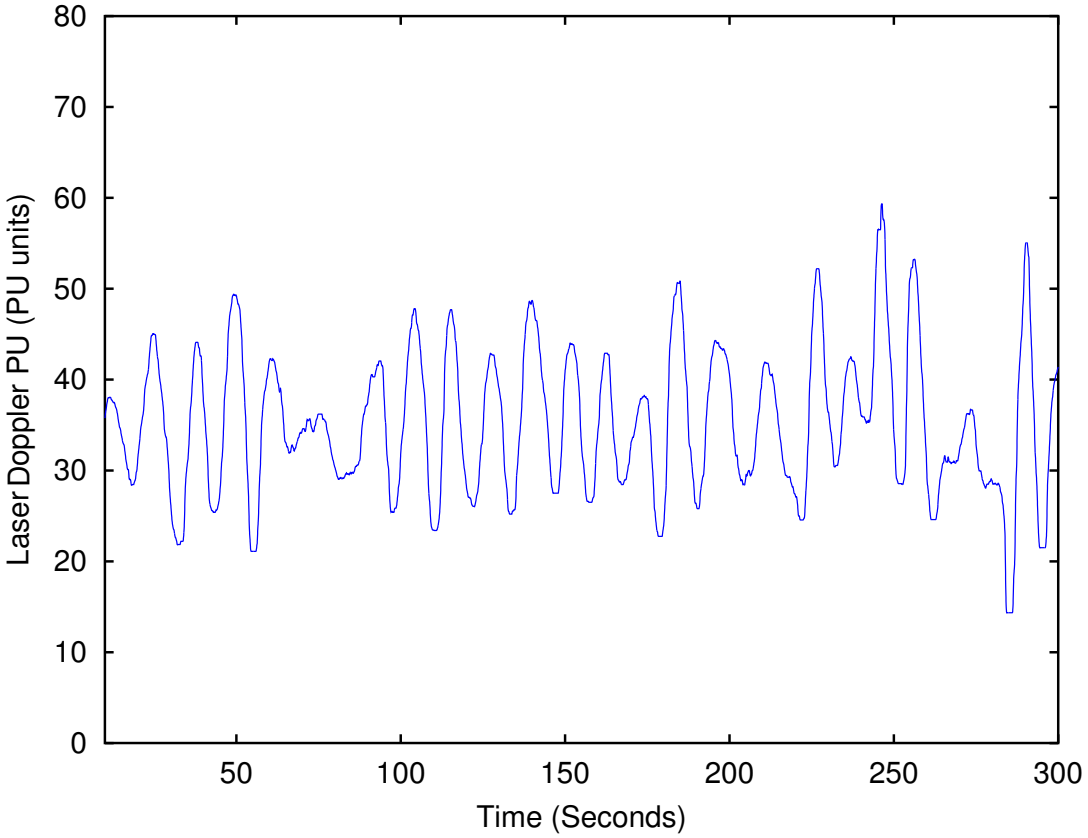


Figure 10.11: Time series filtered laser Doppler data from the Moor instruments device, showing characteristic ~0.1Hz oscillation (V14C data).

From the figure it is clear that there is a large modulation with a frequency centred on 0.09Hz (90mHz). Similar oscillatory behaviour was seen in the data from all 18 volunteers.

Studies by Nilsson et al [117], Johnson [255], and Skarphedisson et al[118] found approximately linear relationships between the perfusion metric data output from laser Doppler devices and true volumetric perfusion as measured by fluid flow in a phantom model, venous occlusion plethysmography, and hydrogen clearance in rats respectively. Thus it would appear that if the oscillatory behaviour seen in the volunteer data is a

real effect, and not an artefact of the data processing carried out on-board the Moor Instruments device, then there is a periodic oscillation in perfusion with a modulation depth of order ±40%. Similar oscillation with comparable amplitude was seen in a study by Bernardi et al [256], where it was described as baroreceptor-sensitive oscillation, arising due to feedback processes regulating blood pressure. This hypothesis of a baroreceptor and sympathetic nervous system origin of the 0.1Hz oscillatory cycle is supported by multiple sources, for example Milliami describes the heart rate and blood pressure fluctuations in the 0.1Hz frequency band as being a marker of sympathetic activity [257]. The term "Mayer waves" is often used to describe this oscillatory cycle, after the description by Sigmund Mayer in 1876 [258].

This oscillation is trivial to remove from the laser Doppler data; for example a second order low pass Bessel filter might be an effective choice, as this filter has a relatively short delay compared to its bandwidth and a flat group delay giving low distortion. However, the perfusion oscillation may have important implications in the use of capillary refill as a perfusion metric. As each capillary refill takes of order 3 seconds, it is effectively a point measurement of perfusion when compared to the 0.1Hz cycle. It thus becomes apparent that if capillary refill time is a function of perfusion, a single capillary refill measurement will only provide an approximate measurement, and that averaging over a considerable number of refills will be required in order to reduce the amplitude of these aliasing effects. Hypothetically it might be possible to somehow synchronise the refills with the 0.1Hz cycle (for example by using analysis of heart rate variability or blood pressure data, as Bernardi et al found that MAP was related to the 0.1Hz oscillation), but it is clear from both figure 10.11 and the published studies that the oscillation is quite erratic, with amplitude and period variations making accurate extrapolation difficult if not impossible.

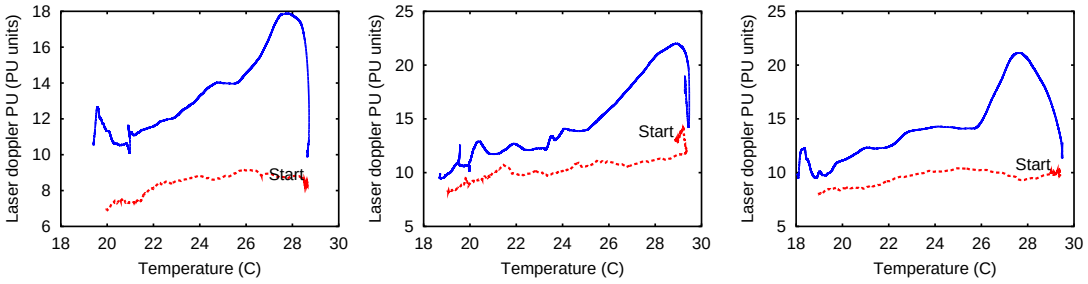
It is interesting to note that published studies of perfusion employing tracer clearance and venous occlusion methods have not shown a 0.1Hz oscillatory cycle. However, it would appear to be likely that these techniques have an intrinsic low pass filter; for a diffusion gaseous tracer such as hydrogen the diffusion dynamics limit the response times, and for venous occlusion plethysmography, the accumulation of blood within the venous compartment may serve to even out perfusion oscillations. An alternative hypothesis might be that changes in the vascular structure occur with a 0.1Hz cycle, and give rise to a oscillation in the laser Doppler PU signal without any perfusion oscillation actually occurring. However, this hypothesis seems unlikely given that a 0.1Hz cycle in cardiac output has been observed using Doppler ultrasound, for example by Elstad et al [259].

To investigate the effect of changing ambient temperature on the laser Doppler PU metric, the raw PU data from the MoorVMS device were passed through the same filtering process as was used to generate figure 10.11, but with the low pass Bessel filter now set to a 2mHz bandwidth. The temperature data from the probe were passed through an identical Bessel filter in an effort to minimize offset and distortion mismatch between the temperature and PU.

Figures 10.12 to 10.24 (solid blue traces) show filtered PU plotted against filtered temperature from the 17 healthy child volunteers and 20 healthy adult volunteers from the PaedCEW study. The red dashed traces show the same raw data processed using a more sophisticated adaptive filtering technique that is discussed in the next subsection of this chapter.

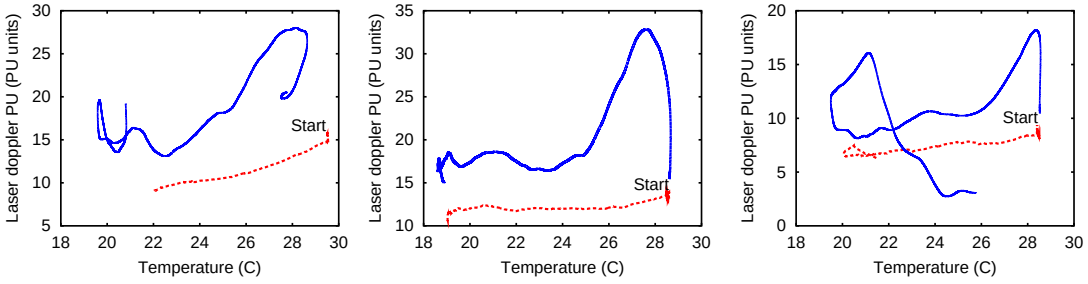
From the blue traces, it can be seen that once the forearm is placed into the chamber and the temperature begins to decrease, there is a general PU trend downwards with decreasing temperature in all but two of the datasets (V7C, and V15C). However, if the baseline PU value before the arm is placed in the chamber is compared to the PU value at the end of the chamber exposure, the pattern becomes somewhat less clear, with V1C, V5C, V6C, V11A, V12A and V17C all having baseline values less than or equal to the value at the end of the chamber exposure. Thus out of a total of 37 volunteers, there was an apparent downward trend in PU once the forearm was placed into the chamber in 95% of the volunteer datasets, and a decrease in PU below the baseline value in 78% of the recordings processed with this relatively simple method. However, the extent to which the presence of the large artefacts within the data obscures the true PU trend is unclear from the analysis.

In adult recordings V1A, V2A, V3A, and V5A, the laser Doppler device was left on the forearm recording data as the forearm warmed following removal from the chamber. It might be expected that a hysteresis loop would be observed upon warming back to room temperature, as skin perfusion returns to the baseline value. Such a response can be seen in the V1A, V3A and V5A recordings, where the upper side of the loop is from the warmup period. However, in the V2A recording there is no such response, and the PU remains relatively stable as the forearm warms. One possible explanation for the absence of any PU increase with rising temperature might be that the experiment was completed before the perfusion started to rise, as it is likely that the physiological response to forearm warming has a certain amount of lag. An alternative explanation is that the laser Doppler sensor changed its orientation and/or its applied pressure on the volunteer's forearm due to the optic fibre lead being repositioned as the arm was placed into and came out of the chamber. It would be necessary to repeat the cooldown/warmup



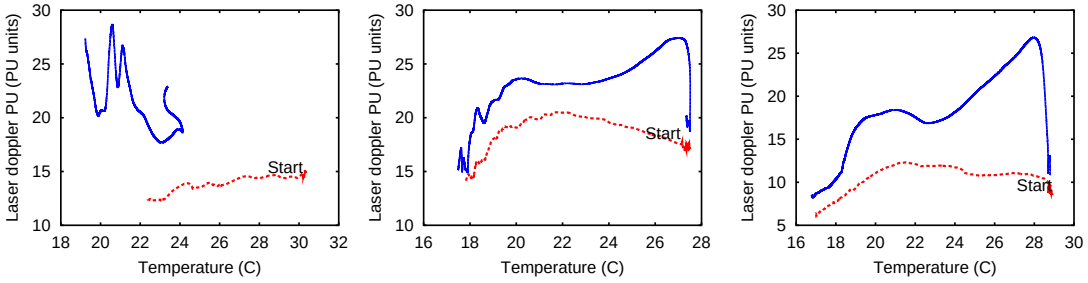
(a) V1C (b) V2C (c) V3C

Figure 10.12: V1C-V3C Laser Doppler



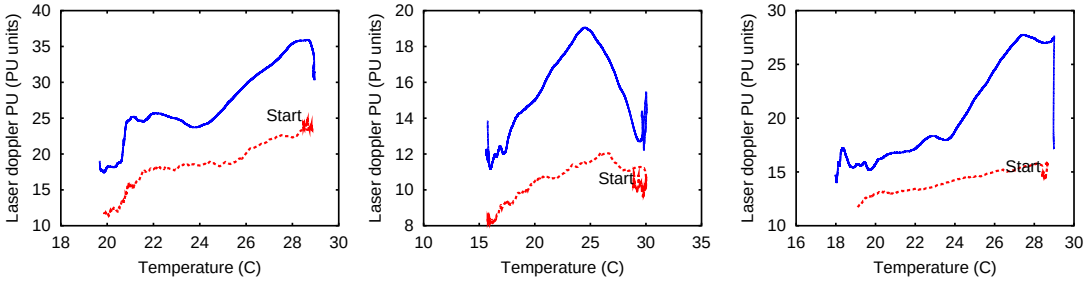
(a) V4C (b) V5C (c) V6C

Figure 10.13: V4C-V6C Laser Doppler



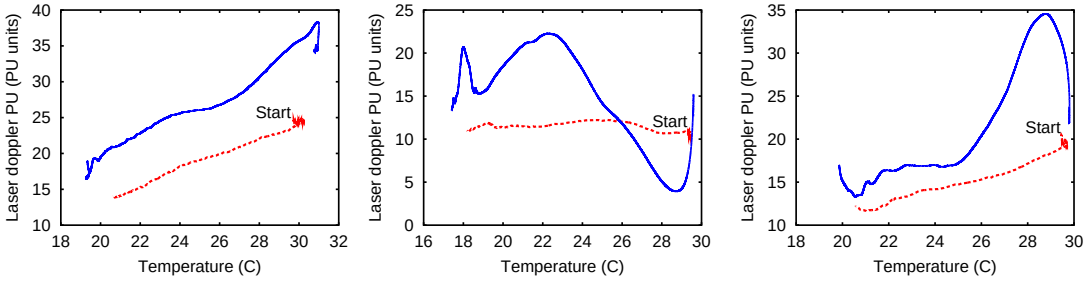
(a) V7C (b) V8C (c) V9C

Figure 10.14: V8C-V9C Laser Doppler



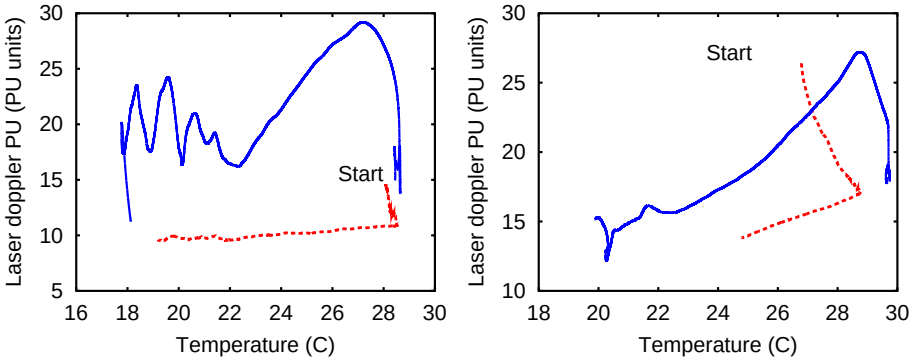
(a) V11C (b) V12C (c) V13C

Figure 10.15: V11C-V13C Laser Doppler



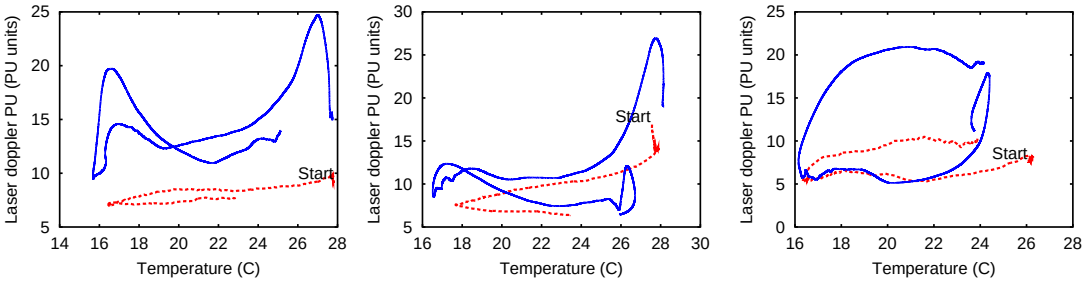
(a) V14C (b) V15C (c) V16C

Figure 10.16: V14C-V16C Laser Doppler



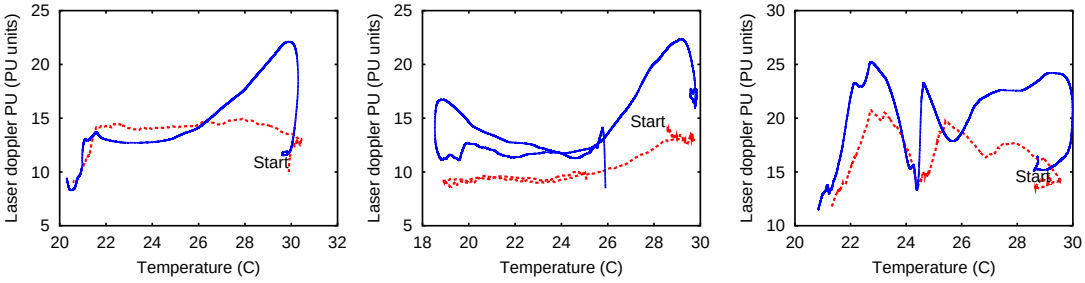
(a) V17C (b) V18C

Figure 10.17: V17C-V18C Laser Doppler



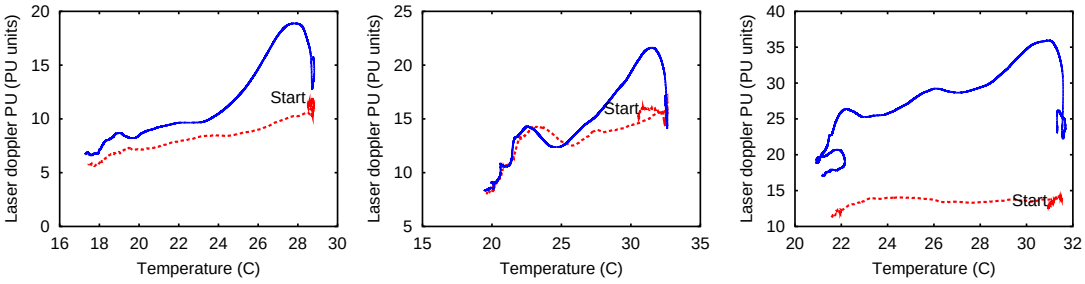
(a) V1A (b) V2A (c) V3A

Figure 10.18: V1A-V3A Laser Doppler



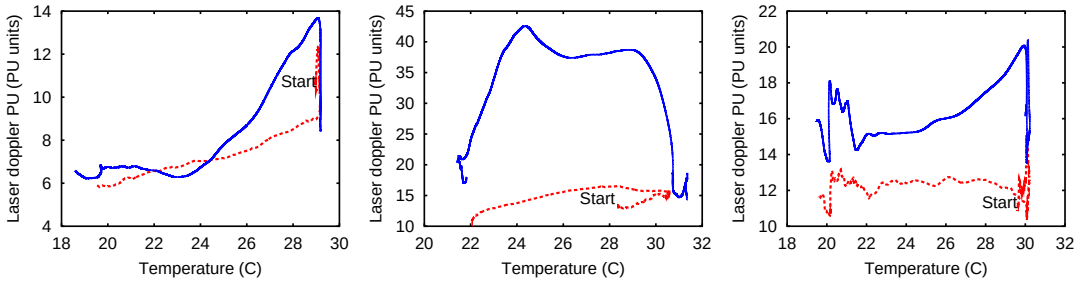
(a) V4A (b) V5A (c) V6A

Figure 10.19: V4A-V6A Laser Doppler



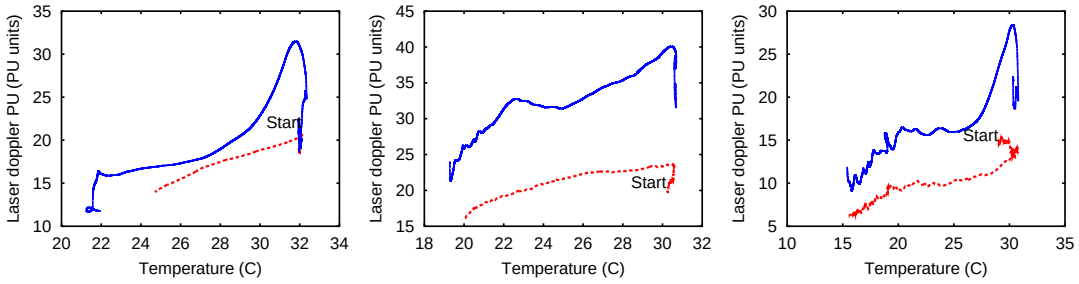
(a) V7A (b) V8A (c) V9A

Figure 10.20: V7A-V9A Laser Doppler



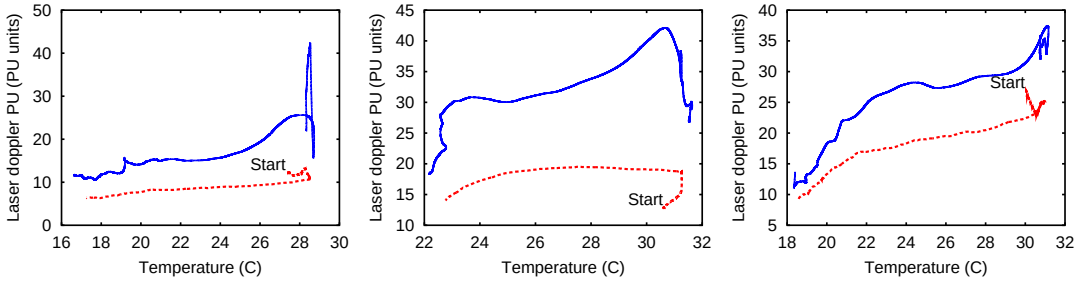
(a) V10A **(b)** V11A **(c)** V12A

Figure 10.21: V10A-V12A Laser Doppler



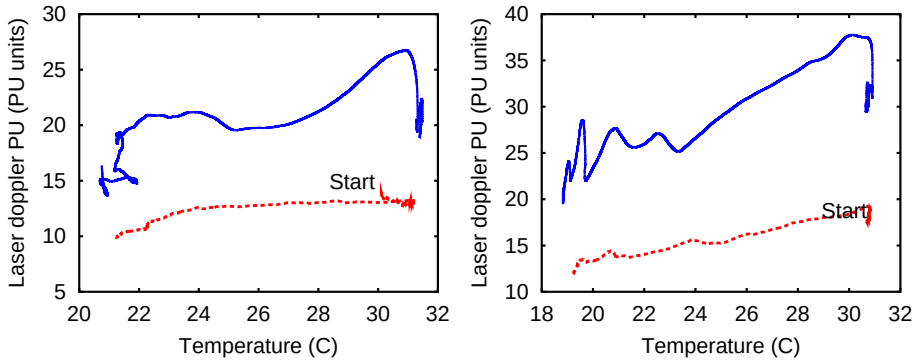
(a) V13A **(b)** V14A **(c)** V15A

Figure 10.22: V13A-V15A Laser Doppler



(a) V16A **(b)** V17A **(c)** V18A

Figure 10.23: V16A-V18A Laser Doppler



(a) V19A **(b)** V20A

Figure 10.24: V19A-V20A Laser Doppler

hysteresis loop experiment with a larger group of volunteers in order ascertain whether sensor orientation changes are significant.

Application of Kalman filtering to MoorVMS data

The raw 40Hz datafile from the MoorVMS device contains PU indices that are unfiltered in the time domain from sample to sample and have no range limit imposed. This clearly represents a large amount of apparently redundant information, as, due to the duration of the experiment, only around 2mHz bandwidth is required for analysis of the Doppler trend with temperature, a reduction of 4 orders of magnitude. This extra information may allow an estimation to be made of the extent to which sensor motion has corrupted the data, and for this to be incorporated as an uncertainty metric into an adaptive filter such as a Kalman filter. However, it is unclear how a Kalman filter could be applied to the raw 40Hz sampled data; the measurement covariance for each sample will clearly be variable, as there are periods of each recording with very significant artefacts, but no intrinsic estimate of covariance is available from the MoorVMS device.

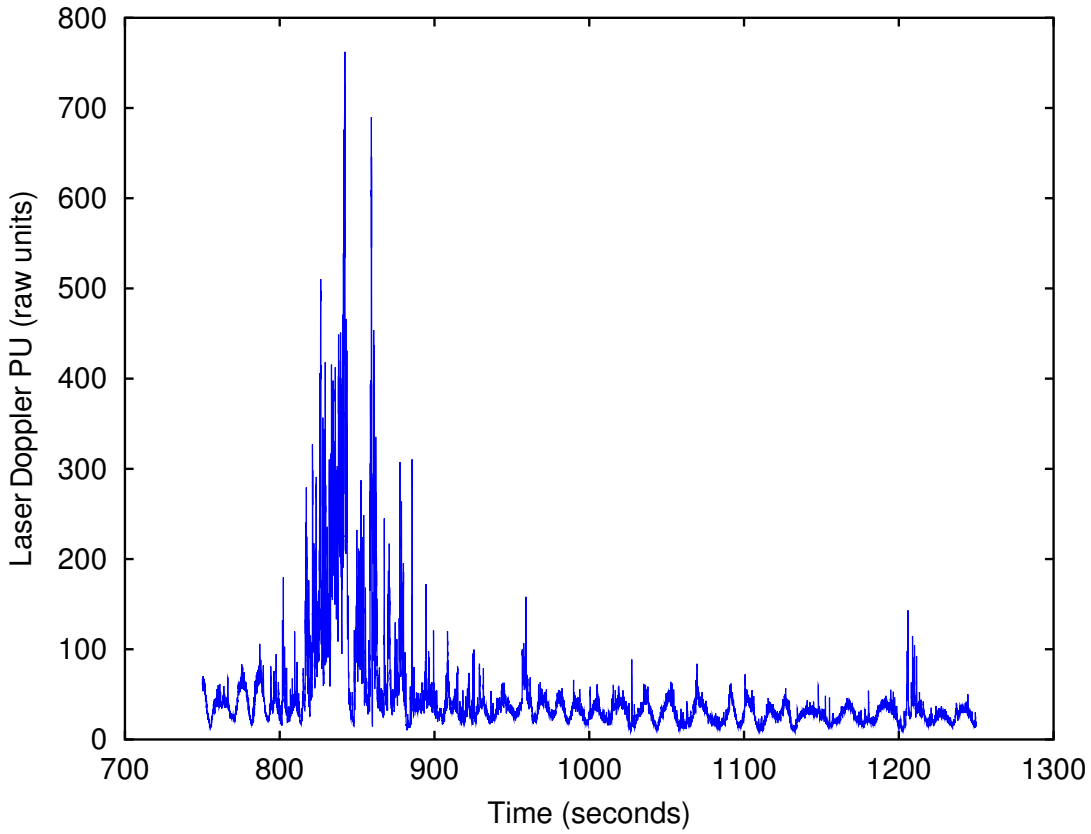


Figure 10.25: Raw unfiltered laser Doppler data from the Moor instruments device, showing characteristic noise spikes (V14C data).

This problem can be resolved if the 40Hz data are processed as a series of windows

(either discrete, sliding or interleaved), with descriptive statistics applied over each window being used to estimate both the "measurement" (of PU), and its covariance for input into a Kalman filter. A Kalman filter is theoretically only an optimal filter if the measurements are subject to Guassian errors, but considering figure 10.25, it can be seen that this will not be the case as the errors are biassed upwards, so a Kalman filter will not be the optimal adaptive filter for the PU data. It may still be a much more effective filter than for example a quantile filter followed by a Bessel low pass, as the Kalman gain can be adjusted in real time to match the fluctuating motion noise. In order for Kalman filtering to be most effective, it is clearly important to choose appropriate functions for estimating the PU measurement and its likely covariance (i.e. the degree of error likely to be present on the measurement), as these need to be at least representative of the true signal and its statistics for the filter to be effective. Considering figure 10.25, it can be seen that the spikes in PU are almost entirely positive, although some negative spiking is seen, for example between 880 and 920 seconds, so the 0.05 quantile was used as a measurement, rather than the minima of the dataset as would be the optimal choice if no negative spiking was present.

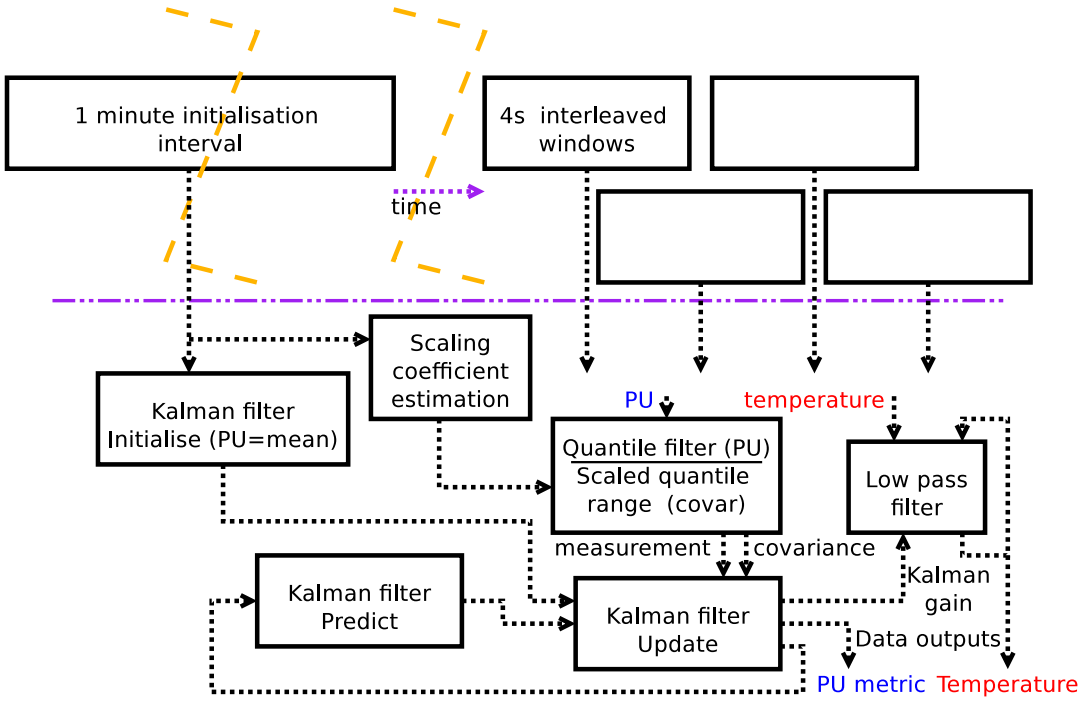


Figure 10.26: Schematic flow diagram of the Kalman filter based laser Doppler processing algorithm.

Figure 10.26 shows a flow diagram of the filtering scheme. As in the previous filtering algorithm, a quantile filter was used to process the raw 40Hz data in each window, but an interleaved rather than rolling windowing was used; before applying the quantile filtering, the data were pre-processed into an interleaved sequence of 4 second long

windows, with the measurement output being the 0.05 quantile of each window interval. The difference between the 0.05 quantile and 0.30 quantile was used as an error metric, based on the observation that the motion spikes are usually of short duration (usually between 1 and 4 samples, i.e. 25 to 100ms), meaning that, over a 4 second window, the presence of spiking causes the spread of the quantile to increase. If there are a small number of spikes, the spread between lower and upper quantiles will increase, but the lower quantiles will not be significantly affected by the presence of spikes, so for this reason the difference between the measurement quantile and the 0.30 quantile was used as opposed to either a higher quantile or the maximum of the window. Theoretically the first derivative of the quantile at the measurement point (0.05) might be a more accurate measurement of this spreading. However, due to the discrete nature of the data in each window, it was found that extending the quantile spread measurement to cover one quartile (0.05 to 0.30) gave a more robust estimator.

To ensure that the filtered PU data could be accurately synchronised with the temperature, a variable gain low pass filter using the same gain as the Kalman filter was applied to the MoorVMS temperature sensor data. This can be seen at the lower right hand side of figure 10.26.

A scaling coefficient was used to convert the squared quantile range into a covariance estimate, with the scaling coefficient estimated for each dataset using the first minute of raw data. This interval of 60 seconds was first divided into 29 interleaved 4 second intervals, with the PU measurement and unscaled covariance estimates then calculated for each window using the process described previously. Assuming that the PU noise process is Brownian, the ratio of the covariance of the first derivative of the PU measurement set to the mean of the set of unscaled covariance estimates can then be used to derive a value for the scaling coefficient. This assumption is consistent with the prediction stage used in the Kalman filter, which was simply a linear addition to the state covariance, i.e. a Lévy process. However, this is clearly not consistent with the data, as it can be seen that there is a roughly linear trend in PU with temperature, and the temperature trend was roughly linear with time for the cooling phase of most recordings. An alternative to this sub optimal design might be a Kalman implementation incorporating a relationship between temperature and PU as its state model, but this would involve an assumption of the model relating PU to temperature, so for this reason it was determined that a random walk type model of PU changes as a function of time was a good choice for the purpose of initial investigation of PU changes with temperature.

From the Kalman filtered traces (red dashed), it can be seen that the laser Doppler

PU value is almost always lower than the quantile and low pass filtered data (solid blue traces), e.g. figures 10.12a and 10.21c, suggesting that interference spikes are significantly affecting the blue traces. As the true perfusion is unknown, it is difficult to assess the quality of the filtered PU signal.

Observing the change in filtered PU over the lower 4°C range of temperature, it can be seen that there is a downward trend in all but six of the 37 recordings (V3A, V5C, V5A, V12A, V17C, and V15C). Alternatively if the "baseline" filtered PU prior to cooling is compared to the PU value at the lowest temperature, all the recordings except V12A, V15C, and V17A show a downward trend. However, the presence of large signal deviations of apparent spurious origin in V6A makes this recording dubious, and the large baseline shift in V18C and V2A also means that the accuracy these recordings is in doubt. For these reasons it would appear that V2A, V6A, V12A, V15C, V17A, and V18C should be categorized as having either an apparent PU increase with cooling or insufficient change in PU compared to the spurious artefacts for a trend to be clear. Thus with the Kalman filtering there was a clear downward PU trend in 31 of the 37 recordings, or 84% of the volunteer datasets.

It is interesting to compare the V1A, V2A, V3A and V5A recordings with the Kalman filtering and quantile/Bessel filtering (red dashed and blue lines in figures 10.18b-10.18c, and 10.19b). In these experiments the MoorVMS device remained attached to the forearm for some time after the arm was removed from the chamber. It can be seen that the cooldown and warmup loop is much tighter in V5A with Kalman filtering, apparently due to a very significant reduction in the effect of spurious motion induced spikes on the filter output. The V1A and V3A loops are also tighter with Kalman, although there is not as dramatic an improvement here. Upon closer inspection of the raw 40Hz data, it was found that the signal artefacts were heavily concentrated around the times when the arm was taken into and out of the chamber, as would be expected if they originated from sensor motion. This effect can be seen in figure 10.25, where the forearm was placed into the chamber starting at 820 seconds. It would thus seem reasonable to conclude that the upward trend in the solid blue filter trace in figure 10.19b and many of the other blue traces upon insertion into the chamber is purely due to the motion artefacts biassing the 0.25 quantile used by the filter.

The reason for the continued PU decrease upon rewarming in recording V2A may have been clarified somewhat by the Kalman filter; there would appear to be a continual decrease in PU with time over the entire duration of this recording, potentially indicating either a real underlying physiological trend in this volunteer, or alternatively a problem with either the sensor or its attachment to the volunteer. As was concluded

previously in the context of the solid blue trace data; with only four volunteer recordings of a cooldown/warmup cycle, it is difficult to make substantive conclusions as to the underlying mechanisms.

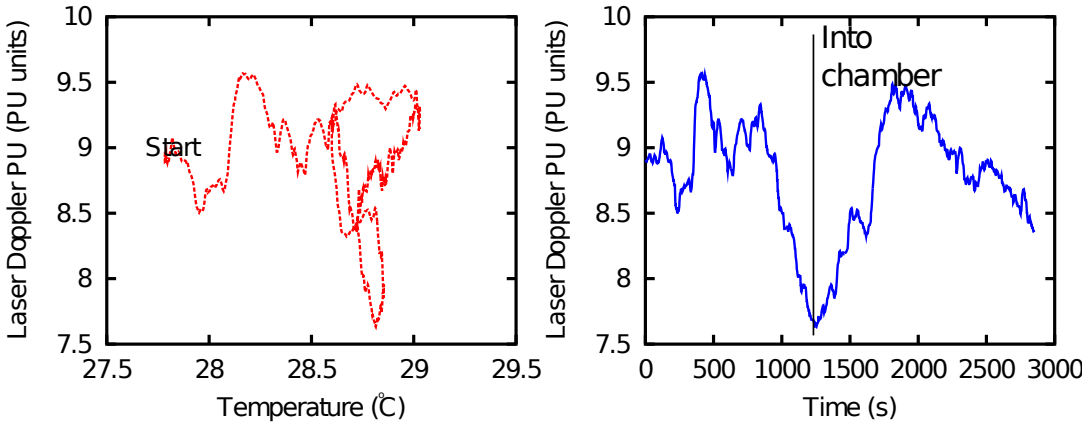


Figure 10.27: Baseline from a healthy adult volunteer showing forearm insertion into an uncooled chamber at 1200 seconds (volunteer 1).

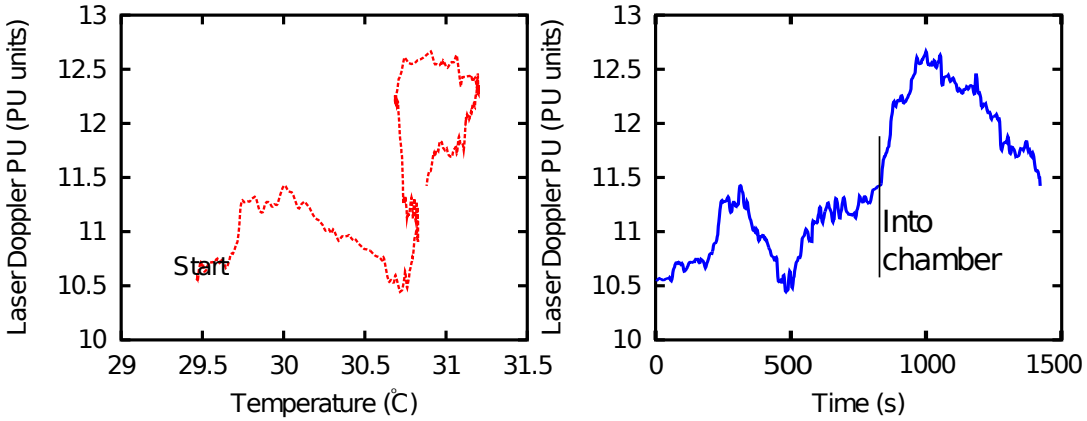


Figure 10.28: Baseline from a healthy adult volunteer showing forearm insertion into an uncooled chamber at 800 seconds (volunteer 2).

Two healthy adult volunteer recordings were made in an attempt to better understand the baseline drift seen in the laser Doppler PU metric. Figures 10.27 and 10.28 show plots of the Kalman filtered laser Doppler PU value versus temperature and time. In these two recordings, the same experimental protocol as in the earlier forearm cooling studies was followed, but the chamber was not turned on, and so remained at close to room temperature. It can be seen that there was an increase in the filtered PU value following insertion into the chamber in both volunteers. This does not appear to be due to motional spikes passing through the filter to influence the signal, as the output then trended down linearly in both cases, rather than exponentially as would be expected due to filter impulse response, and the Kalman gain following insertion into the chamber was sufficiently large that the filter time constant was in the range of 100 seconds with

both datasets.

There is also a pronounced downward baseline drift in the first volunteer beginning at 800 seconds into the recording. However, if the size of the baseline drift is compared to the temperature induced changes seen in the earlier recordings, it can be seen that the mean temperature induced change was of order 5 (although this was highly variable between subjects), whereas the baseline drift was within the ± 1 range, so the temperature induced PU changes are significant compared to the expected baseline drift over the duration of the experiment (at least in these two volunteers).

One potential explanation for all the baseline trend features seen in the two studies is that perfusion drops during the period when the forearm is at rest, as it is not being exercised, with a rise in perfusion following insertion of the arm into the chamber due to the temporary exercise involved in both moving the forearm and repositioning the sensor cabling. A study by Zelis et al using venous occlusion plethysmography suggests that this may be a credible explanation [260], with a significant correlation found between arm exercises and volumetric perfusion. The experimental protocol for the PaedCEW study did not strictly specify that the forearm and hand should be held stationary throughout the experiment, only that arm motion should be kept to a minimum to improve the quality of the data. From both the Zelis results and the baseline data it would seem that differing arm and hand movements between experimental subjects may have been responsible for at least some of the baseline drifts seen in the laser Doppler data.

The 1974 study by Zelis et al [260], did not distinguish between muscle and skin flow, measuring total volumetric flow to the hand. The laser Doppler device in the PaedCEW study only measured skin flow, and thus was not sensitive to muscle flow. An earlier 1964 study by Zelis et al [261] sought to investigate skin versus muscle flow, finding that an initial decrease in forearm skin blood flow on commencement of leg exercise was followed by an increase post exercise, both in the case of strenuous and modest leg exercise. Unfortunately these results are not directly applicable to the PaedCEW study as they involved leg rather than forearm exercises, but in conjunction with the 1974 study they would appear to indicate that skin blood flow it likely to be significantly influenced by even small amounts of arm exercise when compared to the baseline with no exercise.

A further confounding variable in the PaedCEW study may be the degree to which the subjects were exercising prior to the start of the study, and the temperatures they had been exposed to (e.g. sitting in a heated office or car versus walking to the location of the experiment). In the 1974 Zelis study, subjects were allowed to rest for 15 to 30 minutes following attachment of the instrumentation, whereas in the PaedCEW study

the experiment began as soon as the sensors had been attached, although the subjects were instructed to remain seated in the room for 10 minutes prior to commencement of the experiment to acclimatise to the room temperature.

10.5.2 Perfusion Index from Masimo Inc. Radical 7 Pulse oximeter

The Masimo Radical 7 pulse oximeter used for the PaedCEW study stored perfusion index (PI) data to internal memory for the duration of each experiment. These data were later transferred via a USB interface, and it is understood to be resampled but effectively identical to the PI data displayed on the LCD monitor on the front face of the Radical 7 device (the live RS232 output was verified as being the same as the LCD display, but as the internally stored data were difficult to verify categorically due to the device internally delaying and resampling the data).

As previously discussed in the context of apnoea effects, the PI metric output from the Radical 7 pulse oximeter appears to have significant filtering artefacts. Unfortunately, unlike with the Moor Instruments MoorVMS device, no unfiltered or "raw" PI output is available, although an analogue photo-plethysmogram is output from a connector on the rear of the Radical 7. The potential utility of this analogue output has not been researched further, but it may be possible to implement some form of perfusion index by capturing these data to a data logger with an ADC.

As the Radical 7 PI data had a relatively high bandwidth, two traces have been plotted for each recording; one with the unaltered raw data (green dashed), and a second with a second order 2mHz low pass Bessel filter applied (the same filter used for generating the solid blue traces in the earlier MoorVMS plots, and shown here as the blue trace).

Figures 10.29 to 10.32 show the PI data from the PaedCEW study, with each figure showing time series data in the left subplot as two traces. Temperature data were available from the MoorVMS laser Doppler device for all the recordings, so a second subplot showing low pass filtered PI versus temperature was also generated. In the majority of the recordings, the two instruments started recording data at different times, so it was necessary to crop some data to allow PI to be plotted versus temperature. A

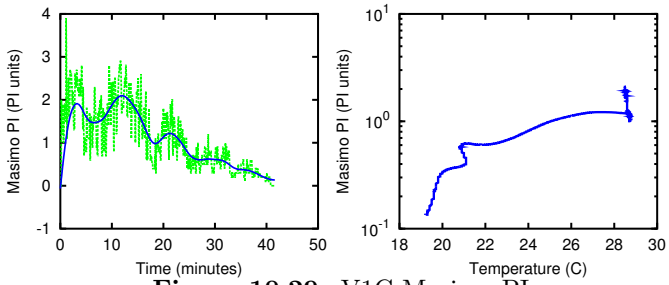
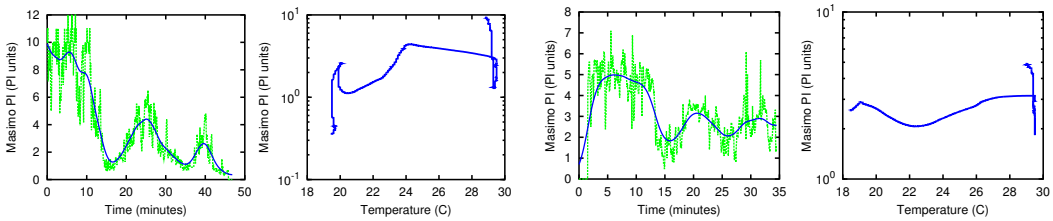
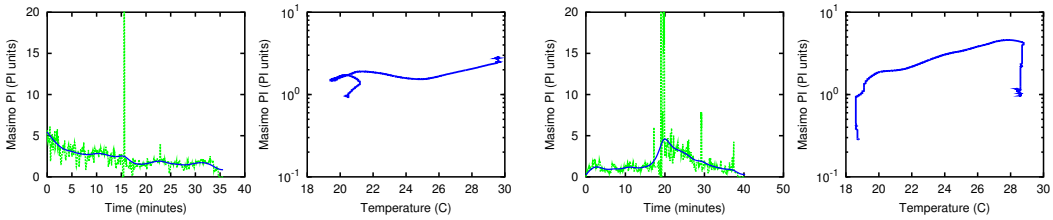


Figure 10.29: V1C Masimo PI



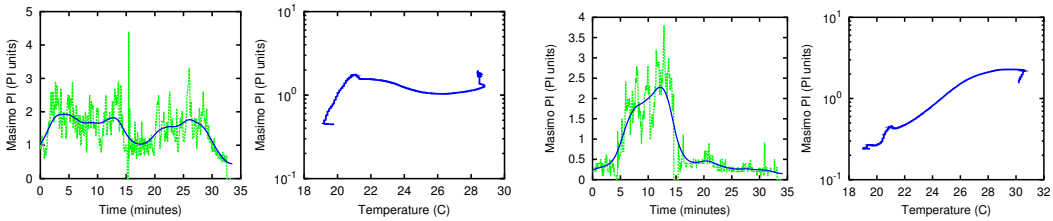
(a) V2C

(b) V3C



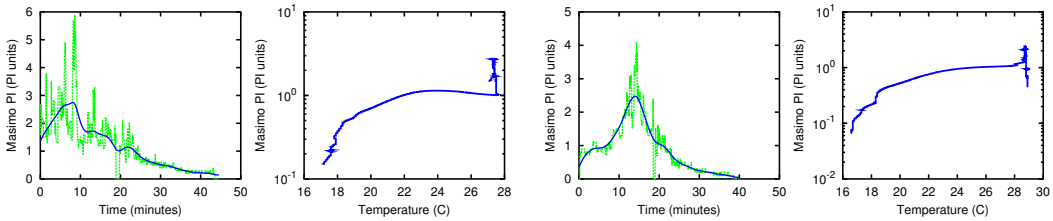
(c) V4C

(d) V5C



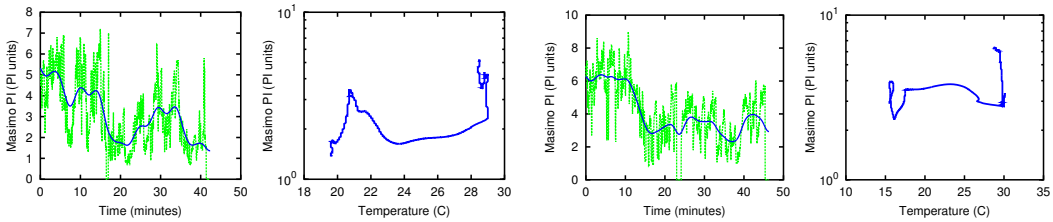
(e) V6C

(f) V7C



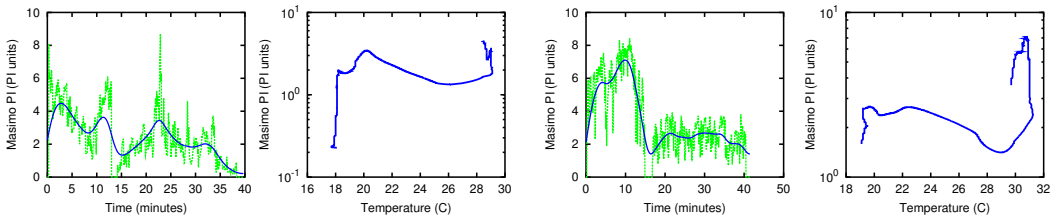
(g) V8C

(h) V9C



(i) V11C

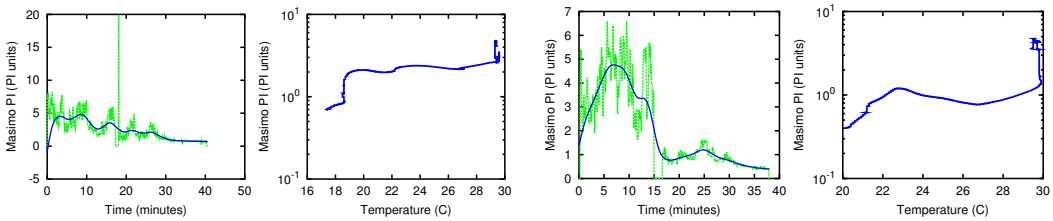
(j) V12C



(k) V13C

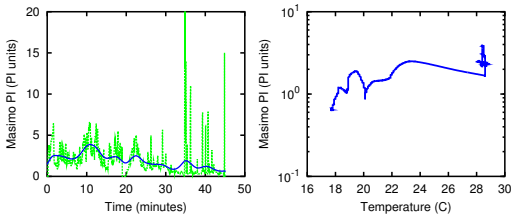
(l) V14C

Figure 10.30: V2C-V14C Masimo PI

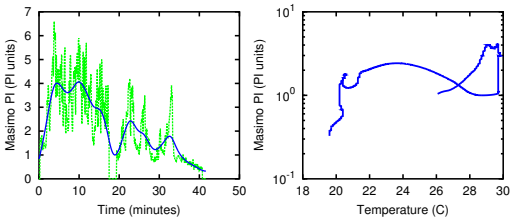


(a) V15C

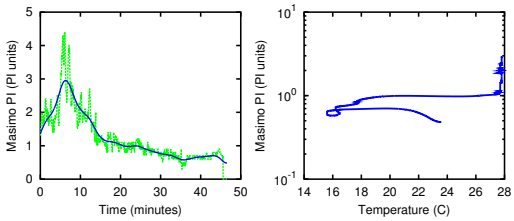
(b) V16C



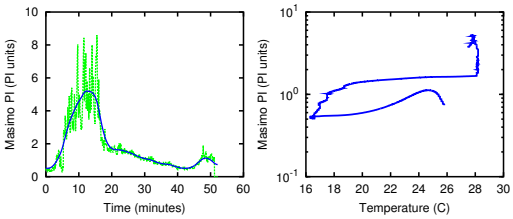
(c) V17C



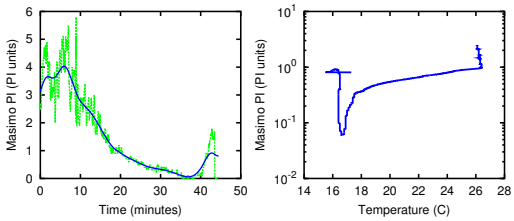
(d) V18C



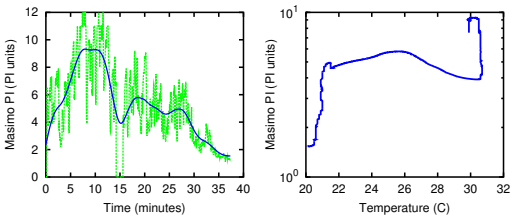
(e) V1A



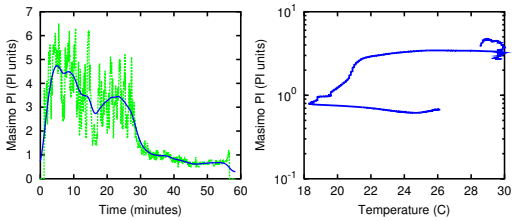
(f) V2A



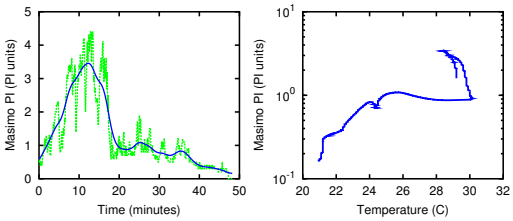
(g) V3A



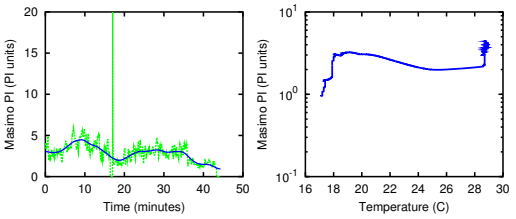
(h) V4A



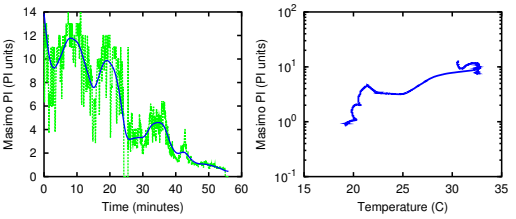
(i) V5A



(j) V6A

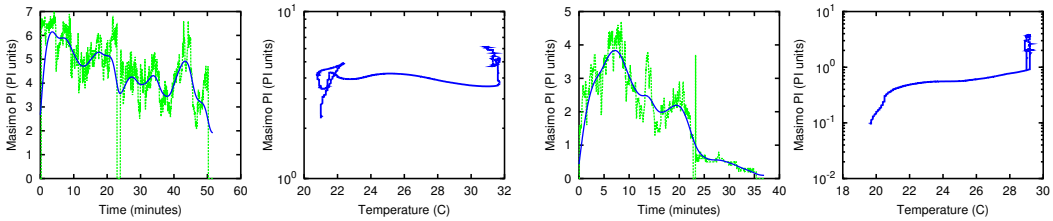


(k) V7A

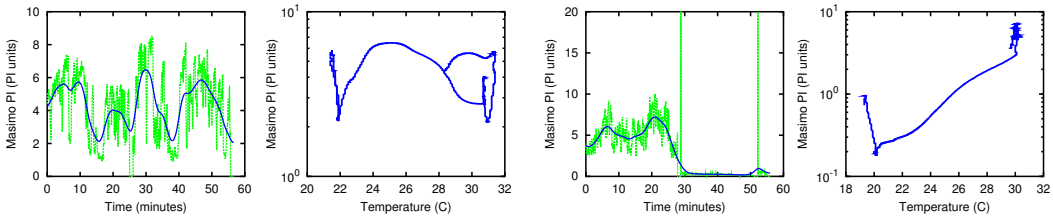


(l) V8A

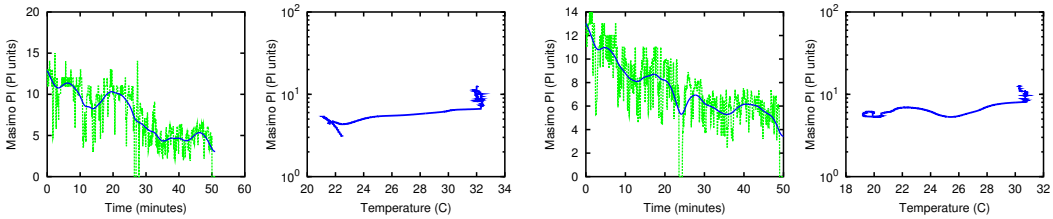
Figure 10.31: V15C-V18C and V1A-V8A Masimo PI



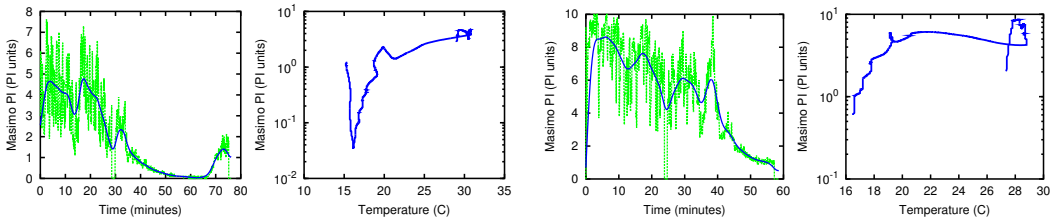
(a) V9A **(b)** V10A



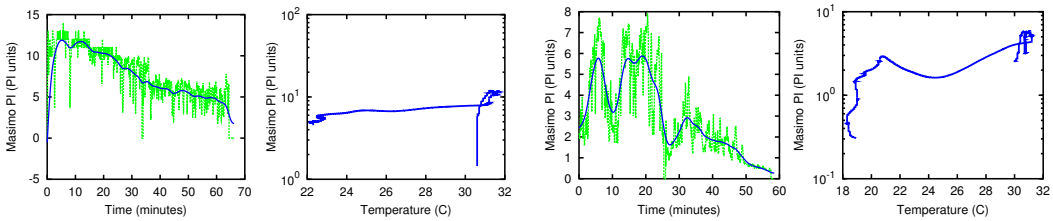
(c) V11A **(d)** V12A



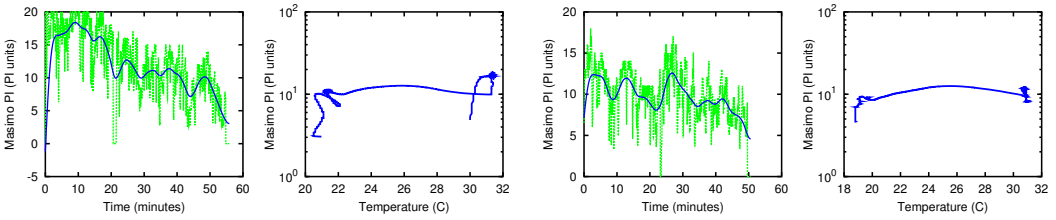
(e) V13A **(f)** V14A



(g) V15A **(h)** V16A



(i) V17A **(j)** V18A



(k) V19A **(l)** V20A

Figure 10.32: V9A-V20A Masimo PI

semi-logscale plot has been used for the PI versus temperature trace, as this was found to give a clearer depiction of the PI range, which dropped close to zero in some recordings.

From the figures it is immediately clear that, in many of the volunteer recordings, there was no clear Radical 7 PI metric decrease below baseline as forearm temperature was reduced. Compared to the Kalman filtered MoorVMS laser Doppler data discussed earlier, it can be seen that the PI metric is relatively ineffective at detecting any clear temperature induced trend.

However, out of the total of 37 recordings, some downward trend can be seen in the majority, with none of the figures showing a clear upward PI trend as temperature decreased. It was noted during the volunteer study that the finger clip pulse oximetry probe tended to become dislodged from the volunteer's finger over the course of the experiment. This led to a loss of signal in many of the recordings, for example at 28 minutes into V12A (figure 10.32d). It appears likely that, in several of the recordings, reattachment of the finger clip resulted in large changes in the PI metric, for example in figure 10.32d the raw PI decreased by an order of magnitude immediately following sensor reattachment. The raw data from adult volunteer V8A and child volunteers V16C, V14C, V13C, and V7C also shows significant changes in PI following sensor dislodgement. In some recordings (e.g. V12A) there are also large positive spikes in the raw PI, suggesting that some signal artefacts are passing through the filtering carried out onboard the Masimo device.

For these reasons it would appear that the apparent decrease in PI seen in recordings V8A, V7C, V13C, V14C, V16C, and V12A arose at least partly due to the change in sensor position on the subject's finger. Of the other recordings, although a majority had a downward trend in PI over the course of the experiment, recordings V4C, V8C, V13A, V14A, V3A, and V9A had an approximately continual decrease with time throughout the entire recording period, i.e. starting before the forearm was placed into the chamber. The data from adult volunteer V5A show a continual decrease in PI after the volunteer removed their arm from the chamber and the temperature began to increase, and so this dataset was also not included in the group of recordings with a significant trend relative to the baseline. Comparing figures 10.31i and 10.19b, it can be seen that both PI and the Kalman filtered laser Doppler data for this volunteer are stable with a slight increase after the forearm was withdrawn from the chamber, as might be expected due to perfusion increasing back to the room temperature baseline level. However, it is interesting to note that the laser Doppler PU metric indicates a continual decrease upon warming in volunteer V2A, whereas the PI indicates an increase back towards the pre cooling baseline.

If a similar success criterion as was applied earlier to the MoorVMS laser Doppler data is applied to the Masimo PI data; i.e. either a baseline range that is above the PI value at the lowest temperature, or a downward trend over the lower 4°C of temperature, then of the remaining recordings that have not already been discarded, only V2A, V4A, V5A, V6A, V1C, V2C, V9C, V15C, V16A, V17A, V17C, V18A, V18C, and V20A show an approximately continual decrease over the lowest 4°C of temperature. This is a total of 14 recordings or 38% of the volunteers. Using the alternative definition of a PI value at the lowest temperature that is below the baseline range, then V1A, V5C, V6C, V7A, V8A, V11C, V15A, and V19A can also be defined as having a significant trend with temperature. This is a total of 22 recordings, or 59% of the results. This compares to figures of 84% for either of these criterion applied to the MoorVMS laser Doppler datasets recorded simultaneously (when Kalman filtering was used to process the raw 40Hz PU output). It should be noted that a more rigorous analysis is conducted in section 10.5.4.

10.5.3 Threshold based capillary refill time with a pneumatic PPG sensor

The capillary refill device used as part of the PaedCEW study was described earlier in section 10.1.2. Due to the choice of optical wavelengths and variation in source to detector separation (i.e. between different LEDs and the photodiode), it was judged not to be feasible to apply NIRS techniques to the processing of the raw intensity data from the experiments. Instead a simple threshold based technique was applied to the data, with a refill time in seconds being estimated for each release of pressure.

Skin surface and capillary refill sensor temperature

The data logger unit described in Chapter 9 recorded the optical power in each of the three wavelength channels, along with the air pressure in the hose to the sensor at a sample rate of approximately 62Hz. Temperature as measured by a thermistor contained within the prototype sensor was also recorded to the data logger in the PaedCEW study. Also as part of the study, a GE Healthcare Dash3000 monitor with attached temperature sensors was used to monitor skin surface temperature every two minutes throughout all the recordings at a site on the forearm adjacent to the capillary refill sensor. The baseline temperature measurements from the pneumatic sensor was found to be between 1°C and 2°C warmer than the Dash3000 temperature probe measurement in all the volunteer recordings.

During the cooling period when the forearm was placed into the chamber, the Dash3000 monitor consistently recorded a larger drop in temperature than the prototype refill sensor, with the Dash3000 temperature decrease being up to double that recorded by the refill sensor. This effect might be explained by the fact that the prototype sensor was relatively large and composed of materials that have low thermal conductivities. Although the Dash3000 sensor has a layer of insulating foam adhered to its top face to reduce the influence of ambient air temperature, the diameter of the sensor (1cm) is much less than the diameter of the prototype refill sensor, which is approximately 4cm. This size difference likely caused the capillary refill sensor to act as an insulating blanket, causing the upwelling body heat flux to cause a temperature rise beneath the sensor, and the heat capacity of the tissue to act as a "thermal mass", reducing the rate of cooling of the tissue under the sensor considerably compared to the adjacent uninsulated skin surface.

To investigate this effect further, a non contact infrared thermometer was used to measure skin surface temperature at one minute intervals.

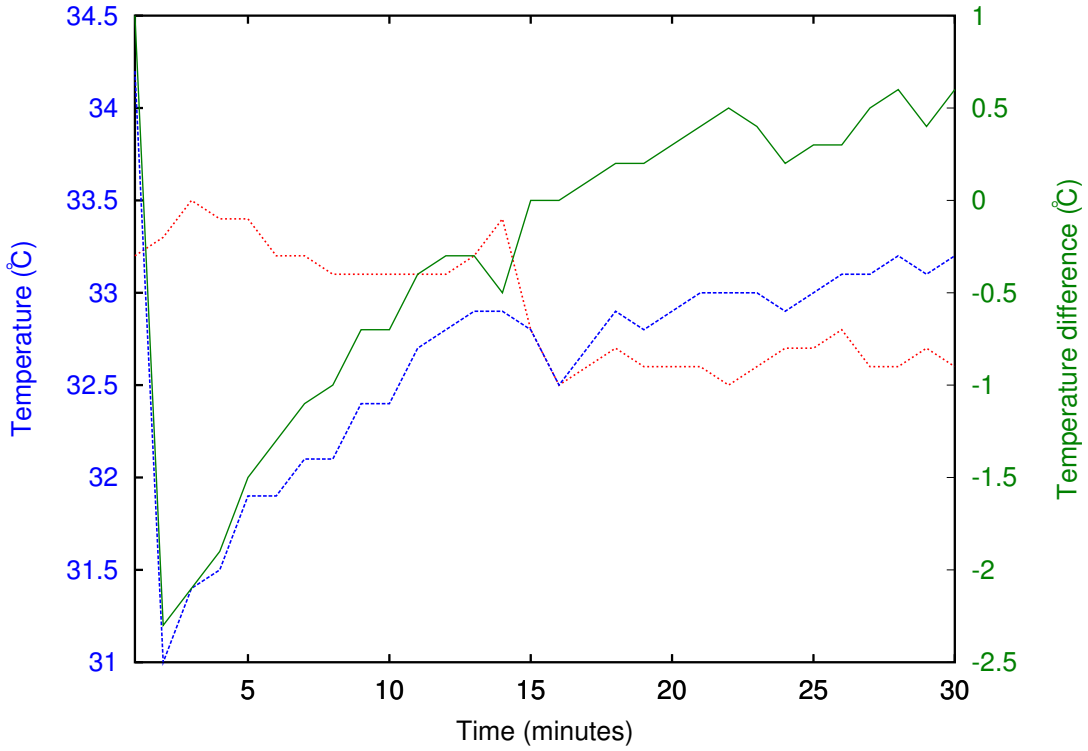


Figure 10.33: Temperature measured beneath the capillary refill sensor (blue dashed) and at an adjacent site (red) during a room temperature baseline study. Arm was deliberately repositioned at 15 minutes.

Following the same experimental protocol as used for the PaedCEW studies, the prototype capillary refill sensor was secured onto the forearm of a volunteer. The outline of the sensor was then marked on the volunteer's arm using a pen, and an adjacent site

(which would normally be the attachment site of the Dash3000 temperature monitor) was also marked. At one minute intervals the sensor was momentarily removed, and the skin surface temperature recorded. From figure 10.33, it can be seen that the sensor was initially at a considerably lower temperature than the skin surface, causing an initial reduction in temperature. This then recovered with an approximately exponential decay, eventually approaching an equilibrium point with the temperature beneath the sensor elevated by approximately 0.7°C compared to the adjacent reference point.

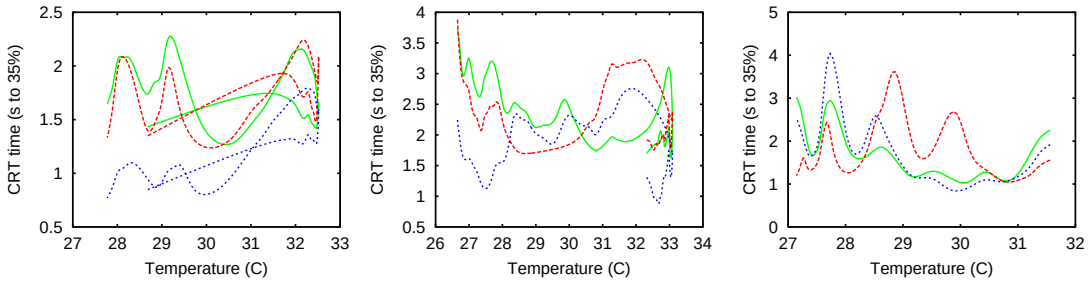
This initial rapid temperature rise explains some of the results seen in the threshold capillary refill time versus temperature plots, for example in the case of volunteer V2A (figure 10.40b) the sensor temperature increased by over 7°C at the start of the experiment. Unfortunately it would appear that neither the sensor storage prior to the experimental recordings nor the procedures for securing the sensors to the volunteers were adequately described in the experimental protocol for the PaedCEW study, resulting in considerable variation in the sensor to skin surface temperature difference at the start of the experiment. This may have had some influence on the inter subject repeatability of the capillary refill time.

Threshold based capillary refill time estimation

A simple threshold technique was used to arrive at a time value for return to a fractional value of the baseline intensity for each capillary refill. This method was applied to each of the three wavelength channels stored by the data logger, leading to three separate refill time traces versus temperature in the following figures.

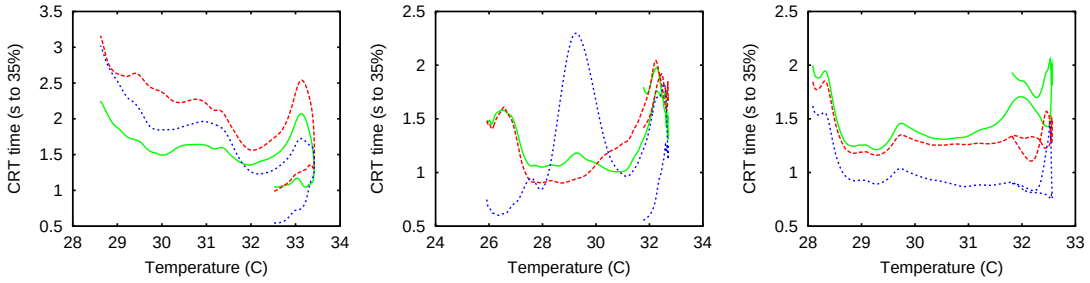
A script written in the GNU-Octave language was used for processing the data, with the air pressure data being used to window the intensity data, isolating each refill interval as the ten seconds following each point where air pressure dropped below a 0.45PSI threshold having previously risen above 0.5PSI and remained above that level for at least one second. A fitting process was then used to produce a baseline and "data quality" for each refill, based on the minimum fit error from windowed line fits to two regions of intensity data, as the transition time between these regions (refill and baseline) was varied from one to six seconds after the pressure release. Refills were excluded if they had an excessive fit error, a positive gradient for the first line fit, or an excessive gradient for the second fit (i.e. too much baseline drift). Rejection thresholds were chosen based on visual inspection of the intensity curves from the unrejected refills, with the threshold being increased to approach the point where clearly erroneous refills were used. A statistical technique may give improved optimisation here. Finally, refill times were calculated for the remaining refills using a threshold method (i.e. time for

intensity to decay to some fraction of the initial height above baseline), before being resampled in the time domain to give a consistent 0.1Hz refill time sample rate. Figures 10.37 to 10.46 show the resulting refill time versus temperature traces, the blue dotted trace is the 950nm channel, the red dashed is 640nm, and the solid green 535nm.



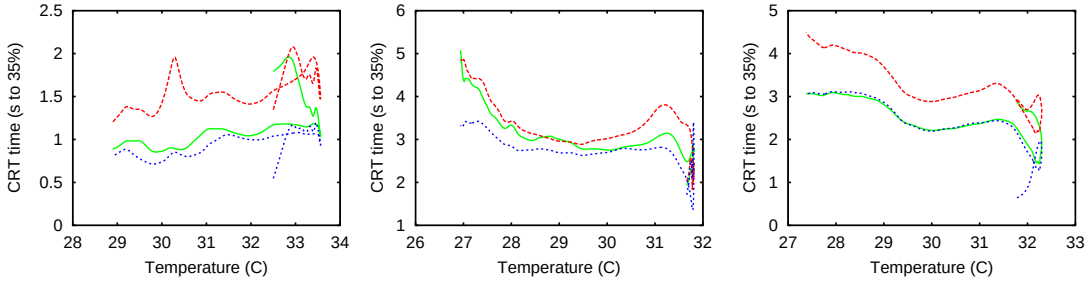
(a) V1C **(b)** V2C **(c)** V3C

Figure 10.34: V1C-V3C Capillary refill time



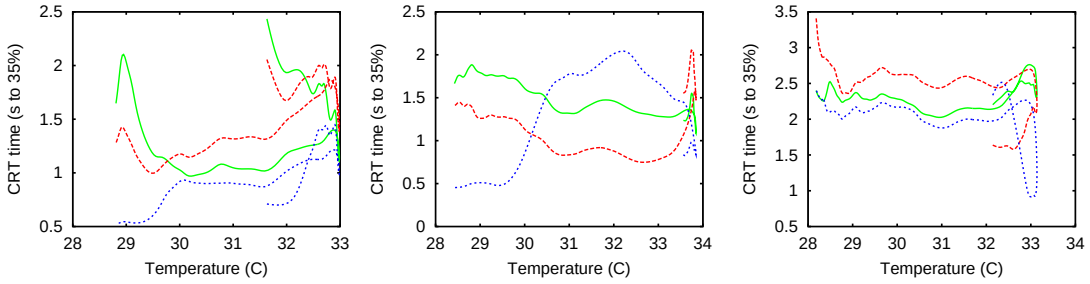
(a) V4C **(b)** V5C **(c)** V6C

Figure 10.35: V4C-V6C Capillary refill time



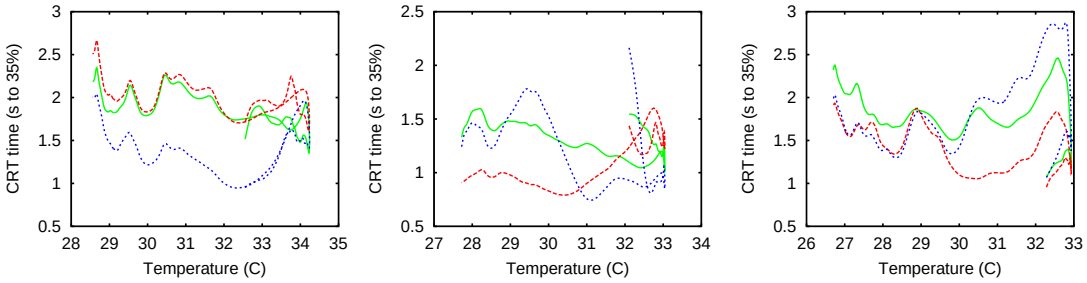
(a) V7C **(b)** V8C **(c)** V9C

Figure 10.36: V7C-V9C Capillary refill time



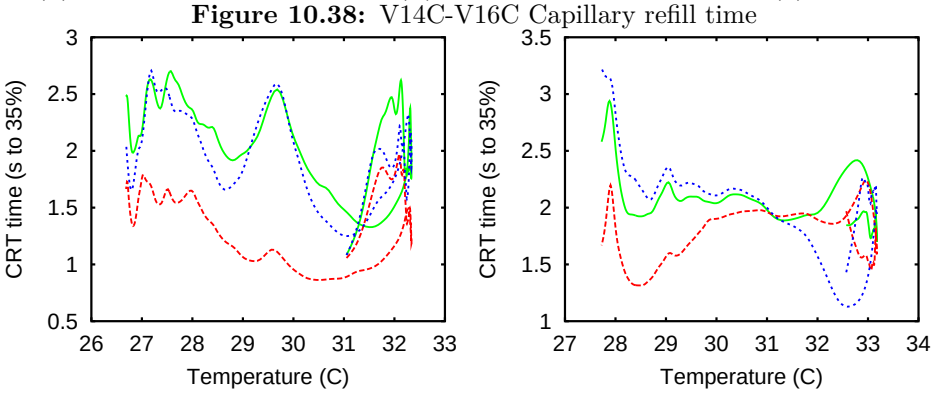
(a) V11C **(b)** V12C **(c)** V13C

Figure 10.37: V11C-V13C Capillary refill time



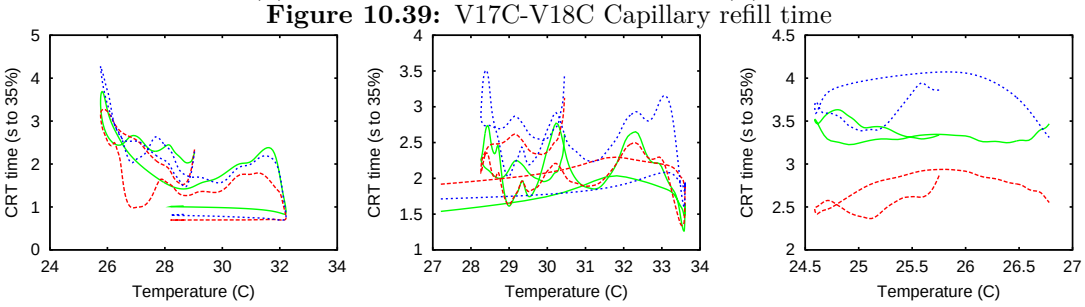
(a) V14C **(b)** V15C **(c)** V16C

Figure 10.38: V14C-V16C Capillary refill time



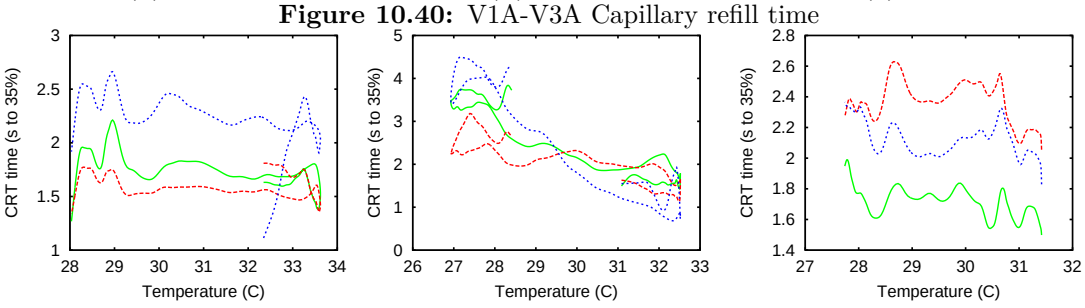
(a) V17C **(b)** V18C

Figure 10.39: V17C-V18C Capillary refill time



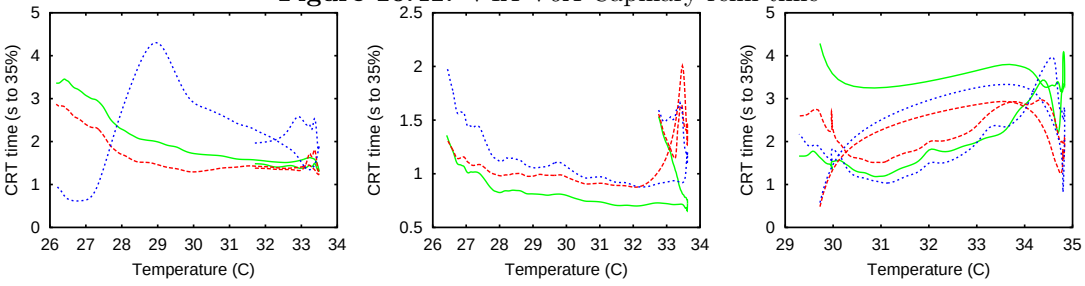
(a) V1A **(b)** V2A **(c)** V3A

Figure 10.40: V1A-V3A Capillary refill time



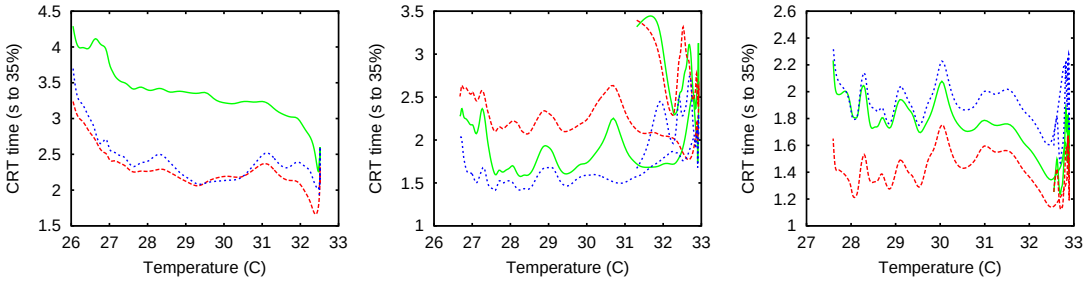
(a) V4A **(b)** V5A **(c)** V6A

Figure 10.41: V4A-V6A Capillary refill time



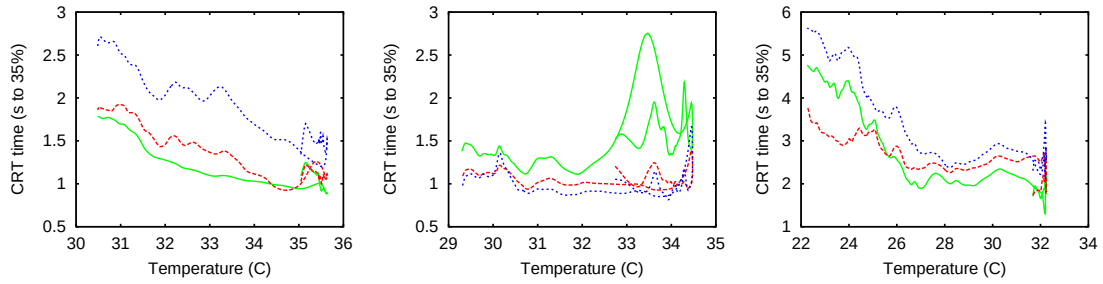
(a) V7A **(b)** V8A **(c)** V9A

Figure 10.42: V7A-V9A Capillary refill time



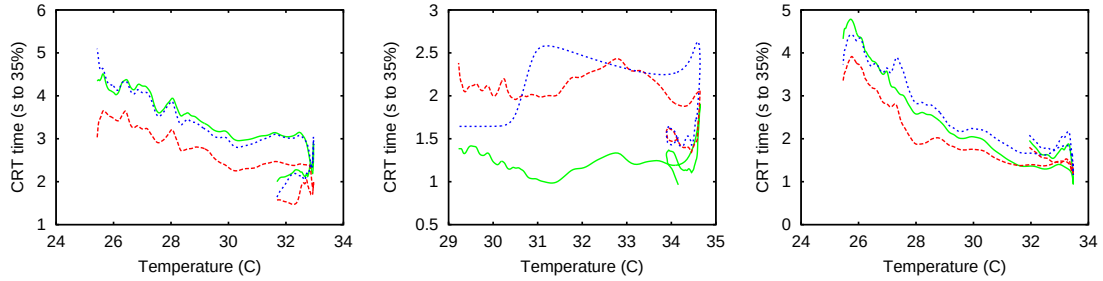
(a) V10A **(b)** V11A **(c)** V12A

Figure 10.43: V10A-V12A Capillary refill time



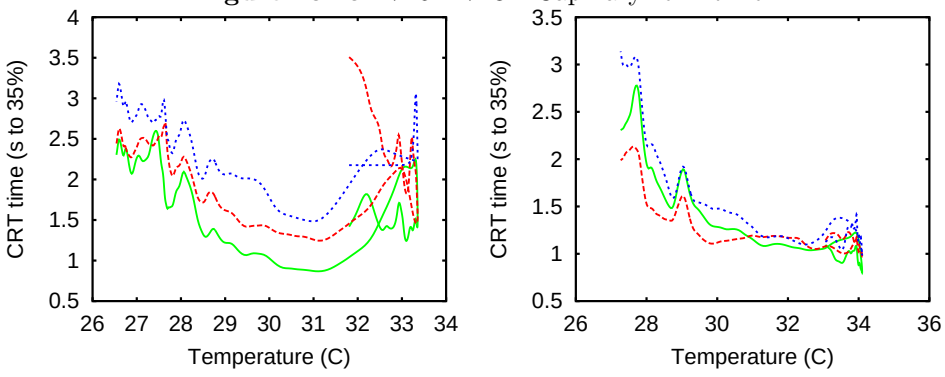
(a) V13A **(b)** V14A **(c)** V15A

Figure 10.44: V13A-V15A Capillary refill time



(a) V16A **(b)** V17A **(c)** V18A

Figure 10.45: V16A-V18A Capillary refill time



(a) V19A **(b)** V20A

Figure 10.46: V19A-V20A Capillary refill time

Prior to plotting, a second order low pass Bessel filter with a bandwidth of 2mHz was applied to both the temperature and refill time. The 35% threshold was based on optimisation of the Pearson r values (described in the following subsection); the threshold value being incremented in 5% steps from 5 to 95% until the "optimal" point, where the greatest percentage of volunteer recordings saw a significant negative correlation between refill time and temperature was found at a 35% threshold. It would appear that a low threshold results in excessive influence of baseline intensity drift and motion artefacts upon the refill time, and that a higher threshold results in a reduced refill time variation with temperature compared to the intrinsic refill timing variability. Performance was not changed very dramatically as the threshold varied; changing the threshold to 35%±10% had only a small effect, with the number of significantly correlated recordings dropping by only 12%.

Pearson r value calculation

Earlier in this chapter, a semi-qualitative approach was used to analyse the PaedCEW study data from the MoorVMS (laser Doppler) and Masimo Radical 7 (pulse oximetry based PI) devices. In the case of the Masimo Radical7, it was not possible to use the Pearson r value as a significance test, due to the unknown nature of the serial correlation introduced by the device. In the case of the MoorVMS device, a simpler raw data output was available, which was processed using a Kalman filter. Unfortunately it is hard to quantify the accuracy of the Kalman filter's state propagation model, so the variance estimate from the Kalman filter state is not an accurate significance test, and due to the variable gain of the filter, serial correlation of the filter state is also difficult to quantify, meaning that a Pearson r significance test would be flawed.

However, as the capillary refill test consists of multiple independent observations separated in time, a correlation test would appear to be applicable. The Pearson product moment coefficient or Pearson r value of two datasets X and Y is defined as follows:

$$P_{X,Y} = \frac{cov(X,Y)}{\sigma_X \sigma_Y} \tag{10.24}$$

if a Fisher transformation is applied, the standard error in the transformed space for a sample size of N is then:

$$SE = \frac{1}{\sqrt{N-3}} \tag{10.25}$$

so, returning to the original untransformed product moment coefficient space, the 95% confidence interval will be

$$r = tanh\left(arctanh(r) \pm \frac{1.96}{\sqrt{N-3}}\right) \tag{10.26}$$

and a significant correlation or anti-correlation can be identified as it will have a confidence interval that does not cross zero. However, for significance testing to a p value of $p < 0.05$, the 90% confidence interval should be used, as, assuming a symmetrical error distribution, there is a 50% change of errors being in the direction that crosses zero r.

This analysis assumes that both datasets (X and Y) consist of completely independent measurements (i.e. the measurements in each set are unrelated, but the two sets may be related to each other). Although it is easy to process each capillary refill time as an independent measurement, there may still be serial correlation that is biological in origin. Earlier analysis of the MoorVMS laser Doppler data indicated significant perfusion oscillation at around 0.1Hz (figure 10.11), which may alias significantly with the capillary refills. Potentially this could be serially anti-correlated, leading to an overestimated error on r. Alternatively, slowly varying bio-regulatory drift will tend to bias the error estimates downwards.

Analysis of the serial independence of the temperature data is somewhat simpler; as the construction details of the sensor are known, it was possible to use the finite element modelling scripts developed for the thermal sensor analysis in Chapter 2 to find the thermal time constant of the sensor (actually the combination of thermal diffusion and thermal mass effects from layer of epoxy resin between the thermistor and the skin surface and the FR4 circuit board behind it), which was approximately 8 seconds in the simulation, meaning temperature data-points are almost completely independent. It should be noted that this is much shorter than the time constant of approximately ten minutes for the entire sensor/skin system, as can be seen from figure 10.33, but that the thermal time constant of the entire system includes the influence of the large thermal mass of the tissue below the sensor, and as the goal of the significance test is to investigate the relation between tissue temperature and tissue blood flow, only the time constant of the sensor itself is of relevance for determining if individual data-points are statistically independent.

The exclusion criterion used in the generation of the refill time versus temperature plots was applied prior to calculating the Pearson r values presented in table 10.6. In the table the data from each volunteer are classified as having a significant or insignificant trend according to whether the confidence interval includes zero. Some of the volunteers had a significant positive correlation, i.e. they saw an increase in refill time with temperature; where a decrease would be expected if capillary refill time was an indicator of perfusion. These are indicated as "(signif.)" on the table.

Volunteer	Pearson r (with confidence interval) for each wavelength channel					
	535nm		640nm		950nm	
V1C	-0.22 +0.25 -0.22	insignif.	0.01 +0.26 -0.27	insignif.	0.44 +0.19 -0.24	(signif.)
V2C	-0.46 +0.16 -0.13	signif.	-0.16 +0.20 -0.19	insignif.	-0.17 +0.20 -0.19	signif.
V3C	-0.50 +0.41 -0.27	signif.	-0.15 +0.38 -0.34	insignif.	-0.40 +0.40 -0.29	signif.
V4C	-0.40 +0.23 -0.19	signif.	-0.54 +0.22 -0.17	signif.	-0.75 +0.16 -0.10	signif.
V5C	0.15 +0.19 -0.20	insignif.	0.36 +0.17 -0.20	(signif.)	0.40 +0.18 -0.22	(signif.)
V6C	0.22 +0.20 -0.22	(signif.)	-0.06 +0.22 -0.22	insignif.	-0.29 +0.23 -0.20	signif.
V7C	0.44 +0.16 -0.20	(signif.)	0.20 +0.21 -0.23	insignif.	0.38 +0.19 -0.23	(signif.)
V8C	-0.63 +0.14 -0.11	signif.	-0.69 +0.12 -0.09	signif.	-0.51 +0.18 -0.14	signif.
V9C	-0.70 +0.12 -0.09	signif.	-0.80 +0.09 -0.06	signif.	-0.82 +0.08 -0.06	signif.
V11C	0.05 +0.22 -0.22	insignif.	0.44 +0.16 -0.20	(signif.)	0.57 +0.14 -0.18	(signif.)
V12C	-0.55 +0.17 -0.13	signif.	0.24 +0.19 -0.21	(signif.)	0.12 +0.26 -0.28	insignif.
V13C	0.10 +0.20 -0.21	insignif.	-0.49 +0.18 -0.14	signif.	-0.49 +0.19 -0.15	signif.
V14C	-0.26 +0.19 -0.17	signif.	-0.22 +0.21 -0.19	signif.	-0.13 +0.22 -0.20	insignif.
V15C	-0.44 +0.19 -0.16	signif.	0.50 +0.15 -0.18	(signif.)	-0.54 +0.18 -0.14	signif.
V16C	-0.26 +0.21 -0.19	signif.	-0.27 +0.20 -0.18	signif.	0.08 +0.21 -0.22	insignif.
V17C	-0.24 +0.19 -0.18	signif.	-0.13 +0.20 -0.19	insignif.	-0.15 +0.22 -0.20	signif.
V18C	-0.26 +0.19 -0.17	signif.	-0.00 +0.20 -0.20	insignif.	-0.33 +0.20 -0.17	signif.
V1A	-0.65 +0.24 -0.15	signif.	-0.59 +0.16 -0.12	signif.	-0.74 +0.13 -0.09	signif.
V2A	-0.20 +0.18 -0.17	signif.	-0.17 +0.18 -0.17	insignif.	-0.35 +0.17 -0.15	signif.
V3A	-0.15 +0.28 -0.26	insignif.	0.28 +0.24 -0.27	(signif.)	0.15 +0.33 -0.37	insignif.
V4A	-0.31 +0.20 -0.18	signif.	-0.17 +0.21 -0.19	signif.	-0.29 +0.21 -0.18	signif.
V5A	-0.76 +0.09 -0.07	signif.	-0.61 +0.13 -0.10	signif.	-0.85 +0.07 -0.05	signif.
V6A	-0.19 +0.24 -0.22	insignif.	-0.11 +0.25 -0.23	insignif.	-0.15 +0.25 -0.23	insignif.
V7A	-0.88 +0.05 -0.04	signif.	-0.65 +0.13 -0.10	signif.	0.51 +0.16 -0.20	(signif.)
V8A	-0.29 +0.16 -0.15	signif.	0.30 +0.15 -0.16	(signif.)	-0.14 +0.17 -0.16	signif.
V9A	0.61 +0.11 -0.15	(signif.)	-0.08 +0.20 -0.19	insignif.	0.40 +0.18 -0.22	(signif.)
V10A	-0.74 +0.11 -0.08	signif.	-0.49 +0.18 -0.15	signif.	-0.33 +0.20 -0.18	signif.
V11A	0.19 +0.16 -0.17	(signif.)	-0.01 +0.17 -0.17	insignif.	0.26 +0.17 -0.18	(signif.)
V12A	-0.25 +0.18 -0.16	signif.	0.02 +0.18 -0.19	insignif.	-0.01 +0.19 -0.19	insignif.
V13A	-0.77 +0.08 -0.06	signif.	-0.67 +0.12 -0.09	signif.	-0.69 +0.11 -0.09	signif.
V14A	0.15 +0.19 -0.20	insignif.	-0.08 +0.18 -0.18	insignif.	-0.08 +0.19 -0.18	insignif.
V15A	-0.78 +0.06 -0.05	signif.	-0.55 +0.11 -0.10	signif.	-0.67 +0.10 -0.08	signif.
V16A	-0.76 +0.08 -0.06	signif.	-0.80 +0.07 -0.05	signif.	-0.74 +0.09 -0.07	signif.
V17A	-0.03 +0.16 -0.16	insignif.	-0.51 +0.13 -0.11	signif.	-0.22 +0.24 -0.21	insignif.
V18A	-0.89 +0.04 -0.03	signif.	-0.86 +0.05 -0.04	signif.	-0.86 +0.05 -0.04	signif.
V19A	-0.27 +0.18 -0.16	signif.	-0.18 +0.18 -0.17	signif.	-0.19 +0.25 -0.23	insignif.
V20A	-0.71 +0.10 -0.08	signif.	-0.55 +0.15 -0.12	signif.	-0.67 +0.13 -0.10	signif.
Percentage significant		68%		49%		57%

Table 10.6: Pearson r values for sensor temperature correlated with threshold based capillary refill time. Results for all volunteers in the PaedCEW study.

The 535nm wavelength has the highest rate of significant negative correlation, at 68%, but this is likely to be related to a firmware bug which affected the red and infrared channels, causing an integer overflow in the DSP processing chain and intermittently inverting the intensity data during some of the refills. The exclusion algorithm identified these events reliably, but the reduced number of refills in the red and infrared datasets resulted in an increase in the confidence interval on r. This can be seen from equation 10.26. It would appear that there is a slightly larger r value in the 535nm channel when the volunteer recordings without a significant negative correlation are excluded. The mean values and confidence intervals for each wavelength channels are;

Wavelength	mean r value of significant volunteers	mean r value of all results
535nm	-0.518 ±0.075	-0.305±0.124
640nm	-0.536 ±0.049	-0.222±0.121
950nm	-0.509 ±0.052	-0.227±0.133

Table 10.7: Summary of capillary refill time product-moment coefficient statistics.

It can be seen from the error estimates (95% confidence interval, using standard error on the mean) that there is no significant difference in mean r between the wavelength channels. This is the case with and without pre-selection of volunteer datasets with a significant correlation (middle column versus right hand column). The r values with pre-selection applied have a much narrower confidence interval on r, but are also much more consistent between the three channels. The variance is approximately halved when pre-selection is applied. It can also be seen that the mean product-moment coefficients in the 640nm and 950nm channels are not as high above the threshold for statistical significance as the 535nm channel.

Unfortunately there were only five volunteer sessions with completely bug free brightness control firmware. As can be seen from table 10.4 and the equations following it, it is not possible to produce an accurate error estimate from such a small sample size, so it is difficult to reach a firm conclusion as to whether there is a significant difference in the performance of capillary refill timing between the wavelengths. However, the data were consistent with the brightness bug being responsible for the poor performance at the two longer wavelengths. The large decrease in variance with pre-selection is strongly suggestive of a subgroup of recordings having suffered from brightness bugs.

If there was a significant effect arising for example due to the shorter penetration depth of green (535nm) light versus the other, longer wavelengths, then a difference in r values between wavelengths amongst the volunteers with a statistically significant negative correlation would be expected, whereas if these recordings have a more significant r due to more functional brightness control, and the brightness control is responsible

for the poor performance in the 640 and 950nm channels, then the converse would be observed. Thus, as the pre-selection of recording sessions with a significant negative product-moment coefficient results in a set of recordings with consistent (overlapping error bounds) mean r values across wavelengths, it would appear that any performance effects related to the choice of optical wavelength are small, and that the intermittent firmware brightness control problems lead to the poor performance of the 640 and 950nm channels, but only during some of the recordings.

Spearman rank correlation was also attempted (raw data not shown for succinctness), but it was found that Pearson r was both higher and more significant (lower p) than Spearman for the majority of the datasets. With Spearman, the percentage of volunteers with a $p < 0.05$ at 535nm decreased from 68% to 62%. The reasons for this difference are not immediately clear, but it may be notable that the relative weighting of the different data-points differs between the Pearson and Spearman correlation coefficients. It might have been expected that Spearman would produce a larger r value, as it allows for a nonlinear relationship between the two datasets, but from the plotted data it can be seen that, although a nonlinear relationship likely exists between temperature and refill time, the "noisy" variability in the refill time is significantly larger than any nonlinear component with temperature. The differing response of Spearman to this high refill time variance would appear to be a likely explanation for the higher values produced by Pearson r.

Correlation between MoorVMS and threshold based capillary refill time

The capillary refill time and Kalman filtered laser Doppler PU metric were compared at baseline versus the lowest temperature state, with percentage change in each metric being analysed due to the high inter subject variability in absolute values. For this analysis, both baseline and low temperature were defined as five minute intervals, with the baseline being the five minutes just prior to insertion into the chamber, and the cooled state being the final five minutes for chamber cooling.

The mean of the Kalman filter output and the individual refill times over these five minute intervals were then found for each volunteer dataset, and used to plot the points in figures 10.47 and 10.48. The error bars on each capillary refill time datapoint were calculated using the standard error on the mean, and show the 95% confidence interval. The error bars for 535nm and 950nm are offset in the horizontal axis of the plots by -1% and +1% respectively for clarity. It should be noted that this analysis assumes each capillary refill time measurement is statistically independent. In the following subsection a relatively weak short period serial correlation is found in some of the datasets, but

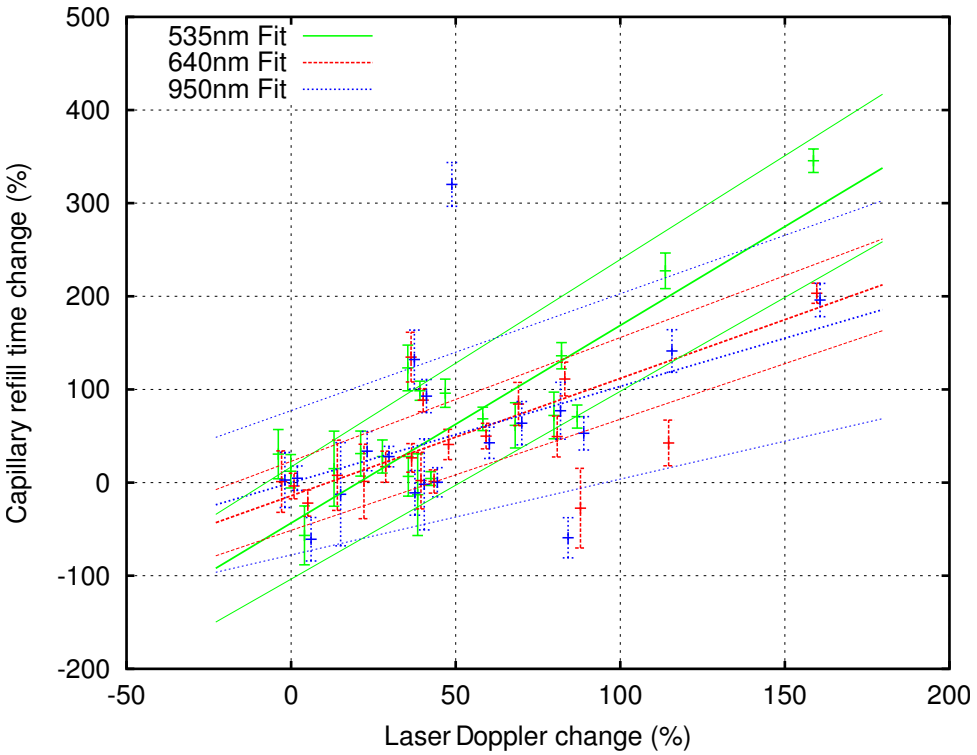


Figure 10.47: Adult volunteers

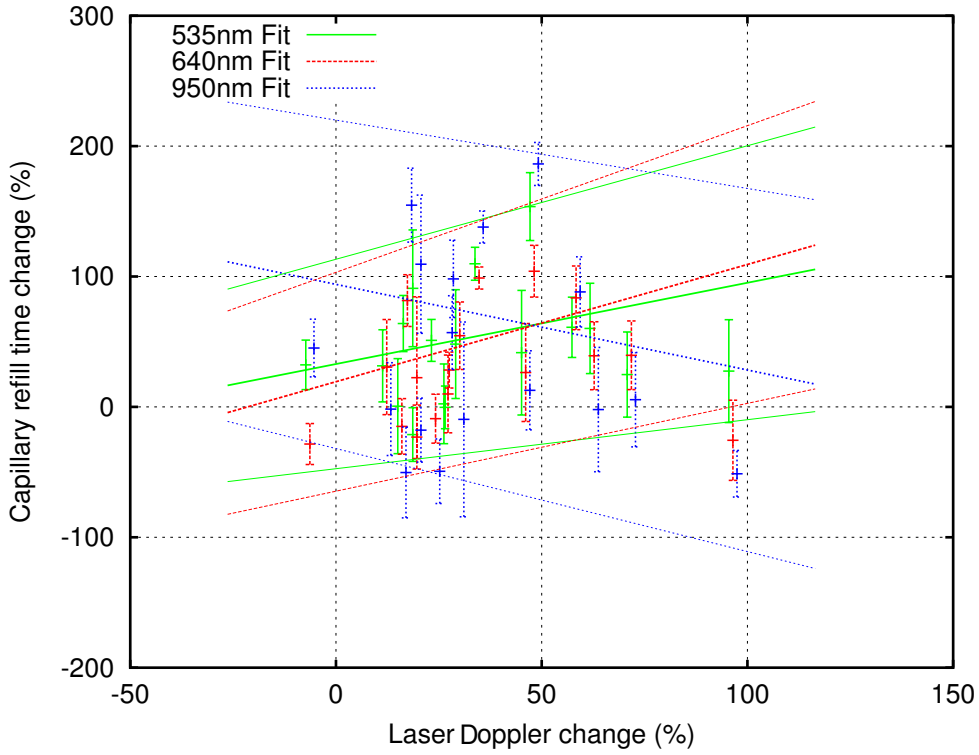


Figure 10.48: Child volunteers

the peak Pearson r value is less than 0.5, so it would appear that the assumption of independence not a gross approximation.

A weighted linear regression is also plotted from each wavelength (solid line), together with its 95% confidence interval (thinner lines). It is immediately clear that the adult volunteers saw a considerably clearer correlation between laser Doppler and capillary refill time change. It is notable that the origin lies with the confidence interval of all linear regressions, meaning that there is no evidence of a systematic trend in one of the metrics that is not present in the other.

Volunteer group	Wavelength	R^2	p
Adults	535nm	0.73	3.2e-6
	640nm	0.40	0.0035
	950nm	0.22	0.04
Children	535nm	0.02	0.58
	640nm	0.03	0.50
	950nm	0.02	0.58

Table 10.8: Summary of correlation data for linear regression fit to the MoorVMS versus refill time change.

Table 10.8 shows the correlation data from the linear regression fits. It is evident that child volunteer datasets have no significant correlation between laser Doppler PU change and refill time change with cooling, whilst there is a significant ($p < 0.05$) correlation at all optical wavelengths in the adult volunteers, with an extremely low p value at 535nm. This wavelength dependence may be related to the firmware problems that caused reduced performance in the 640nm and 950nm channels in some of the recordings, but this does not explain the poor correlation seen in the child volunteers (all under 16 years old). Although it is feasible that there are differences in the cooling response between adults and children, there would appear to be other, more feasible explanations. It may be notable that some of the adult recordings were of a longer duration than the experimental protocol stated, resulting in increased cooling of the forearm (median ΔT=5.5°C, IQR 1.1°C for child volunteers, ΔT=6.3°C, IQR 1.8°C between the five minute averaging windows for adults). It was also noted that some of the younger child volunteers did not stay still throughout the experiment, leading to a decrease in the data quality of both the laser Doppler and the capillary refill data due to motion artefacts.

A useful observation from the adult volunteer linear regressions is the larger percentage change in capillary refill time than in laser Doppler PU (although it is a less statistically significant change), with the percentage change in 535nm capillary refill time being

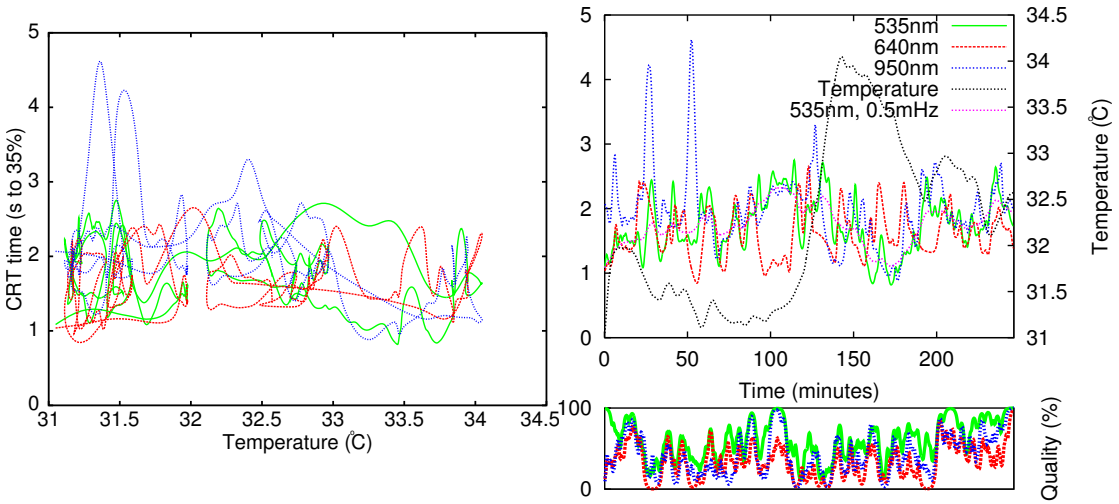


Figure 10.49: Capillary refill time over a four hour period, mauve dotted trace is 0.5mHz lowpass filtered. Quality is unrejected refill percentage.

approximately double that in PU for adult volunteers. This finding may be useful for the scaling of capillary refill time data to laser Doppler PU equivalent or visa versa in clinical monitoring applications.

Baseline analysis of threshold capillary refill time

To determine if there were long (hours) or medium (tens of minutes) term drifts in the capillary refill time, a four hour baseline was recorded with the sensor mounted on the forearm of one of the healthy adult volunteers. During the recording session the volunteer left the recording room and building with instructions to return four hours later. The outside air temperature led to a slight decrease in sensor temperature at the start of the recording, followed by a large rise 2 hours into the recording when the subject entered a warm office environment.

From figure 10.49 it can be seen that there is a considerable range (3°C) in sensor temperature. This compares to a median value of ΔT=5.5°C, IQR 1.1°C for the child volunteers and ΔT=6.3°C, IQR 1.8°C for the adults during the chamber study. A larger value of ΔT=9.9°C, IQR 2.3°C for the entire group was recorded from the temperature sensor on the laser Doppler device, suggesting that the insulating effect of the sensor discussed earlier in subsection 10.5.3 led to a reduction in cooling inside the chamber as well as a warm-up transient.

The refill time from the 535nm and 950nm channels correlated significantly with temperature (r=-0.22 +0.10 -0.09 and r=-0.28 ±0.11 respectively), but no significant correlation was seen between temperature and the 640nm refill time (r=0.03 ±0.13).

In summary; the 535nm channel performed well in the baseline study, with the 950nm channel being slightly worse. Although the 950nm channel refill time was more signif-

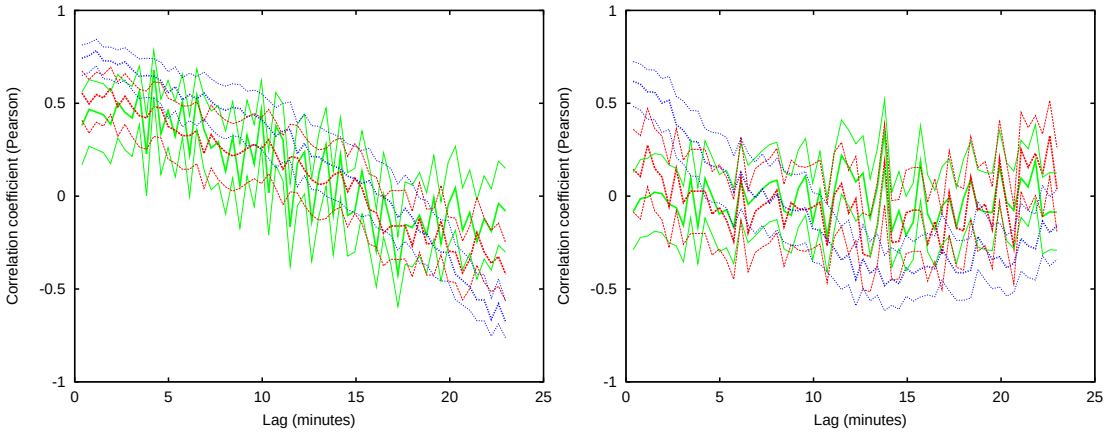
icantly correlated with temperature, there are at least two large anomalous spikes in the 950nm trace during the first hour. These occurred whilst the volunteer was walking with the sensor secured to their arm with tape, and may be related to motion of the sensor (a high percentage of refills were rejected during these intervals).

Time series analysis of threshold capillary refill time

From figure 10.11 it can be seen that a significant oscillation in blood flow is detected by the laser Doppler device. This was discussed earlier in section 10.5.1. By applying time series analysis techniques to the threshold based capillary refill time data, it may be possible to identify this oscillatory cycle in the refill time data. The volunteer datasets "V14C" (healthy child volunteer 14, figure 10.38a) and "V13A" (health adult volunteer 13, figure 10.44a) were selected for this further analysis, based on the observations that V14C had some of the most significant 0.1Hz oscillation in the laser Doppler data, but did not see any significant change in capillary refill time over the course of the chamber cooling, whilst V13A saw some of the most significant refill time changes, but had relatively weak oscillation in the 0.1Hz band (not plotted). Unfortunately it is not trivial to search for a 0.1Hz cycle in the capillary refill data. The capillary refills were made every 23 seconds in the PaedCEW study, much longer than the five second Nyquist period of the blood flow oscillation. However, as the oscillation in the 0.1Hz band appears to have a relatively consistent period, and the interval between capillary refills was fixed at exactly 23 seconds over all the recordings, there may be some aliasing effects.

The Pearson autocorrelation of the threshold capillary refill time was computed as a function of lag time. The rejection algorithm developed for the temperature correlation analysis was again employed here, with rejected refill times being replaced by the "NA" or "Not Available" special number in the GNU-Octave language. This enabled a set of refill times to be constructed with uniform 23 second periods between samples, meaning that a simple indexing offset was all that was required for the autocorrelation function to be computed (although a modified implementation was necessary in order to ignore the NA values). The number of refill times (i.e. number of samples N) was also computed for each wavelength channel, and used to calculate the confidence intervals by application of equation 10.26.

From figure 10.50, a significant long term trend can be seen in the refill time autocorrelation. This related to the change in refill time with temperature, which occurs over a timescale of approximatively 20 minutes in V13A. In the V14C sub figure there is a significant autocorrelation only in the 950nm infrared channel, and from figure 10.38a



(a) V13A **(b)** V14C

Figure 10.50: Refill time autocorrelation function for 535nm (green), 640nm (red dashed), and 950nm (blue dotted). Confidence intervals are shown as thin lines above and below the central thicker line. Data from V13A and V14C volunteers.

it can be seen that there is more variation in the infrared channel over the course of the experiment than in the red and green channels, although this variation was earlier found to not significantly correlate with the temperature.

No short term correlation effects are evident in either subfigure. If there was an influence on capillary refill time from the 0.1Hz oscillation seen in volunteer V14C with the MoorVMS device, then it might be expected that the aliasing of the 0.1Hz cycles would lead to some form of oscillation in the refill time autocorrelation function over the first one to two minutes of lag. Based on the fact that multiple studies have established a direct relationship between true volumetric perfusion of blood and the laser Doppler PU metric, the absence of any such oscillation in the autocorrelation function might then be taken as indicative of capillary refill time not being of value as an indicator of perfusion.

However, considering the oscillatory behaviour seen in figure 10.11, it is evident that the period of the oscillation is variable, and that there are regular jumps in its phase. For this reason it was decided to model the effects of subsampling on the MoorVMS data and compute the autocorrelation function of these subsampled data. Considering a hypothetical situation where capillary refill time was inversely proportional to the laser Doppler PU metric, it is also evident that single point subsampling is not an accurate model. It is clear that capillary refill time cannot be modelled as point sampling of perfusion, due to the means by which the threshold capillary refill time is computed by measurement of a decay over several seconds to some pre-set fraction of the baseline intensity. A full analysis of the relative accuracy of different sampling models would be likely to involve further speculative assumptions, so to avoid this a simple windowed

average of the filtered PU metric over two second periods every 23 seconds was used. The choice of two seconds is somewhat arbitrary, but the mean capillary refill time to 35% threshold was in the region of two seconds, and the refill time was very close to two seconds in the green and red channels throughout the V14C recording.

Prior to resampling, the raw 40Hz sample rate data from the MoorVMS device were filtered using the same process used to generate figure 10.11; exclusion of excessively high and low values, followed by a 1Hz low pass Bessel filter.

The Lomb-Scargle periodogram (originally described Scargle et al [262]) was plotted together with significance levels (α) for the threshold capillary refill time sequence together with this simulated MoorVMS based "refill time" (which was simply the inverse of the filtered PU metric).

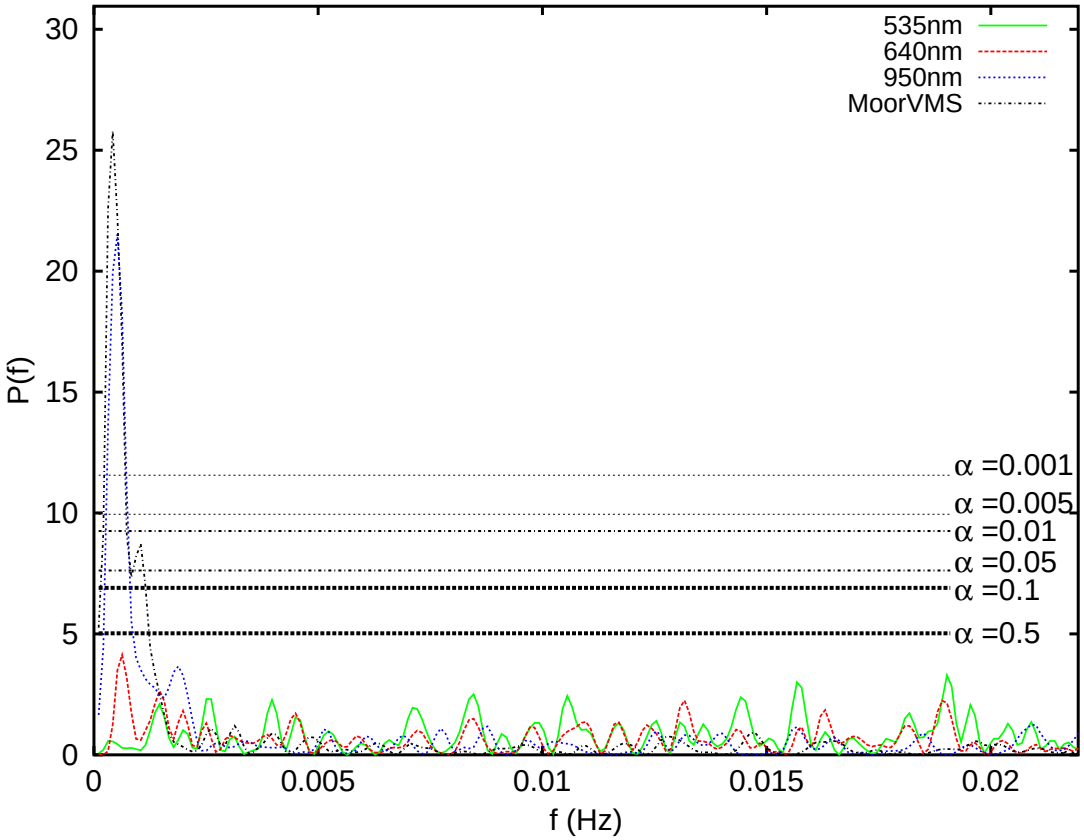


Figure 10.51: Lomb-Scargle periodogram of the threshold capillary refill time data from volunteer V14C, together with the simulated MoorVMS based refill times.

From figure 10.51 it can be seen that the only significant peaks occur at low frequencies in the 950nm capillary refill time trace and MoorVMS simulated trace. These peaks would appear to originate from the long term trend present in the MoorVMS (which has a significant decrease with cooling), and the 950nm refill time. Although some

rippling can be seen at high frequencies in all four traces, there are no coincident peaks between all four, and all peaks are below even the lowest confidence interval of $\alpha = 0.5$. It might be expected that the blood flow oscillation would produce spiking in the 0.01Hz to 0.02Hz band due to aliasing, but it is clear that any such effects are well below the threshold for statistical significance.

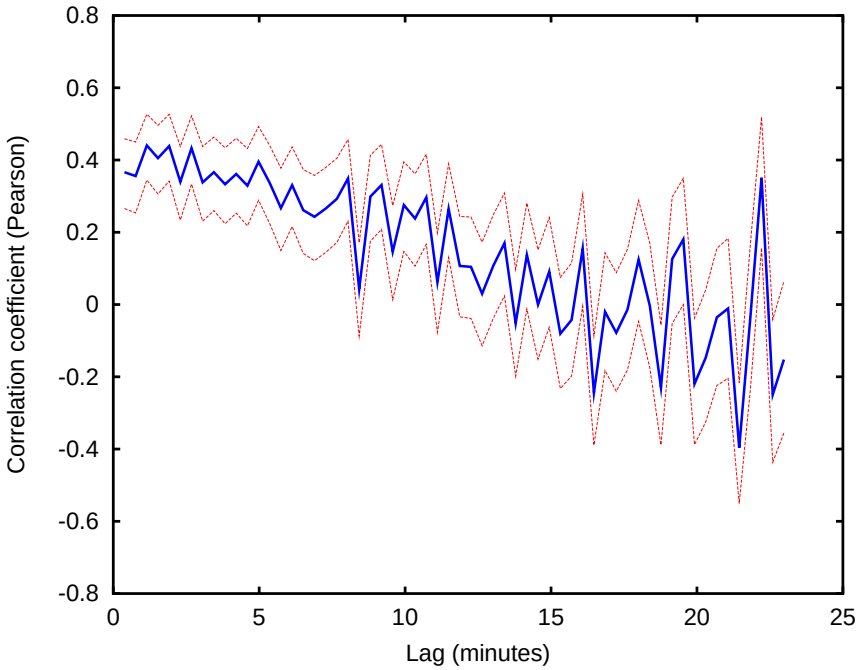


Figure 10.52: Autocorrelation function of simulated capillary refill time. Red dotted line shows confidence interval.

Figure 10.52 plots the autocorrelation function of the simulated refill time. There are also no significant short period features visible here. Only a long autocorrelation decay is present, displaying very similar features to the refill time autocorrelation, meaning that it is not possible to draw any significant conclusions as to the accuracy of capillary refill time by looking for short period autocorrelation features, as these would not be expected to be seen in the autocorrelation function of the entire dataset.

However, upon closer inspection of the filtered MoorVMS data from volunteer V14C, it was found that the period of the blood flow oscillation changed considerably as the forearm temperature was lowered in the chamber. For this reason it was decided to repeat the analysis using only the baseline datasets, i.e. only the data recorded prior to the point at with the arm was placed into the chamber.

Figure 10.53 plots a comparison between the autocorrelation function of the threshold capillary refill time and the simulated refill time generated from the laser Doppler data recorded over the same time period. The confidence intervals are shown with dotted lines, and in the first subfigure the different wavelength channels are indicated by the

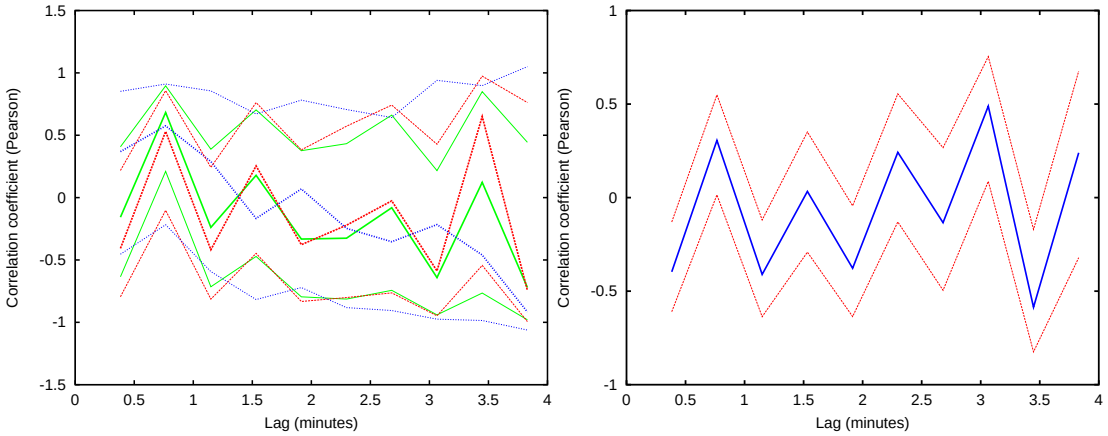


Figure 10.53: Comparison of autocorrelation of capillary refill time sequence before insertion into the chamber with simulated laser Doppler based refill time over this period.

different traces as in the previous figures (green=535nm, red=640nm, blue=950nm). Due to the short duration of the baseline data (7 minutes), only 16 refill times were recorded in each channel, meaning the confidence intervals are too large for any statistically significant conclusions to be drawn for the traces (except possibly the green trace, which has a lower confidence interval bound above zero at a single lag value of 46 seconds). However, oscillation can be seen in the autocorrelation function at all three wavelengths with a very similar period to that seen in the simulated data. The green capillary refill time trace is particularly similar to the simulated trace over the first two minutes of lag.

Although the cropping necessary to produce short period autocorrelation features means the data are not long enough to be statistically significant, the presence of these very similar features in both the simulated and the measured capillary refill time is strongly suggestive of a connection between perfusion and capillary refill time.

If however the simulated refill times generated using the cropped section of baseline data from the MoorVMS device are correlated directly with the capillary refill times recorded over the same time interval, a clear and significant spike is evident at zero offset in all three wavelength channels (figure 10.54). Although a spike always appears at a lag of zero, it was found that the peak correlation only becomes statistically significant over a small two second range of MoorVMS lag time.

To generate figure 10.54, the MoorVMS lag time was first varied over the range -11.5 to +11.5 seconds (i.e. half of the 23 second interval between capillary refills), so that the point at which the mean correlation coefficient across the three wavelength channels was maximised could be found. The figure then shows correlation traces as the lag changes by integer numbers of refills (i.e integer multiples of 23 seconds). The lag interval on

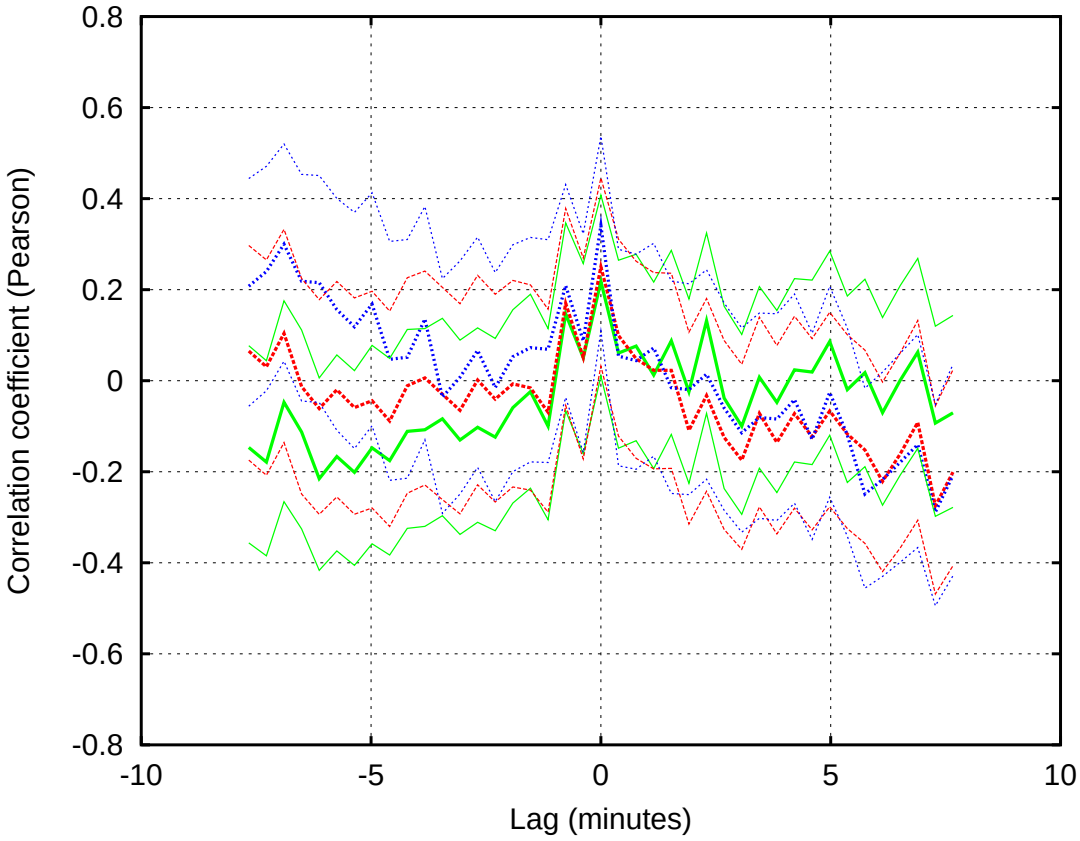


Figure 10.54: Correlation function of simulated capillary refill time with actual threshold refill times. The dotted lines show confidence intervals. The 535nm, 640nm, and 950nm channels are indicated using identical trace styles to previous figures.

the horizontal axis is thus the MoorVMS lag time, i.e. the time in minutes by which the simulated MoorVMS based capillary refill time has been delayed relative to the threshold capillary refill time from the pneumatic capillary refill hardware.

A secondary correlation spike is visible at -0.76 minutes (-46 seconds), as well as a long term drift in the correlation, particularly in the infrared channel. The presence of a long term drift in the correlation between the infrared refill time and MoorVMS would be expected due to the long period drift in both of these datasets. Secondary peaks around zero lag would also be expected due to the nature of the autocorrelation functions. However, only a single secondary peak is seen. It might be expected that the correlation would be approximately symmetrical around zero lag, but the asymmetry indicates that the laser Doppler dataset leads the capillary refill to a more significant degree than the converse relationship. It is unclear why such a relationship would be seen, but it should be noted that the asymmetry is considerably smaller than the confidence interval, and thus is not statistically significant.

It can also be seen that the peak Pearson r correlation coefficient is only 0.36 in the

950nm channel, 0.27 in the 640nm, and 0.23 in the 535nm. These values are not very high (perfect correlation would have $r = 1$), but they are statistically significant according to equation 10.26. Thus in conclusion; the correlation between the threshold capillary refill time and MoorVMS laser Doppler based simulated refill time shows a statistically significant relationship exists between laser Doppler measured perfusion and the capillary refill. The presence of a very similar oscillatory signature in the autocorrelation functions during the baseline at room temperature, as well as the narrow width of the correlation spike at zero lag, further indicates that this correlation arises due to the influence of the 0.1Hz band oscillatory cycle (or Mayer waves) on perfusion.

As multiple published studies have established that the laser Doppler PU is linearly related to true volumetric perfusion, it would seem reasonable to conclude that there is a statistically significant correlation between perfusion and refill time. However, the low correlation coefficient seen in figure 10.54, together with the poor consistency of capillary refill time compared to laser Doppler in the PaedCEW study results indicate that refill time is likely to be a poor perfusion indicator. Furthermore, it should be noted that the statistically significant correlation does not itself demonstrate a causal link between perfusion and capillary refill time.

10.5.4 Summary of PaedCEW study data

The following table (10.9) summarises the data from the preceding subsections; the semi-qualitative analyses of trend over the lowest 4°C of cooling, the baseline levels before and after cooling, and finally the percentage of recordings with significant refill time correlation with temperature.

A modified form of the Newey-West method was discussed earlier in this chapter (section 10.2.1) for estimation of errors in venous occlusion NIRS data. Application of this algorithm to the MoorVMS PU (Kalman filtered) and Masimo PI (raw) datasets was attempted, with the temperatures and PU or PI being converted to deviations relative to their mean, and then zero padded in similar fashion to figure 10.6. An ordinary least squares (OLS) regression was then used to fit the temperature deviation onto the PU or PI deviation, and the confidence interval of the OLS fit computed via Newey-West, with a significant result defined as a confidence interval that did not cross zero. This method was found to massively overestimate error on some of the Kalman filtered MoorVMS data, apparently as the smoothness of the filtered PU dataset caused long periodicity features to dominate the autocorrelation function. A lag limit of 200 seconds was found to prevent this, but is a somewhat ad-hoc solution, so the data shown in the table should be treated with caution. It would appear that the Newey-West method is still a

relatively poor match to the Kalman filtered data, most likely due to the time variant autocorrelation introduced by the variable Kalman gain. Time variant autocorrelation is not modelled in the Newey-West method.

Application of Granger causality [263] was attempted as an alternative means of analysing the Kalman filtered MoorVMS data. Unfortunately this was found to be numerically unstable when applied to the chamber cooling datasets, with the p values diverging to infeasible numbers. This phenomenon is apparently due to the non-stationary nature of the chamber cooling datasets, as application of Granger causality to the baseline study data shown in figure 10.49, (which could be approximated as a stationary process) yields a feasible p value ($p < 0.01$ for temperature Granger causes PU) that is consistent under resampling and windowing of the original dataset.

Metric	Wavelength	Clear trend over lowest 4°C of cooling	No overlap in baselines	Significant correlation
Refill time Adults	535nm	70%	65%	70%
	640nm	55%	55%	60%
	950nm	60%	55%	55%
Refill time Children	535nm	82%	24%	65%
	640nm	65%	24%	35%
	950nm	53%	35%	59%
MoorVMS Laser Doppler Adults	N.A.	85%	80%	80%*
MoorVMS Laser Doppler Children	N.A.	82%	88%	82%*
Masimo PI Adults	N.A.	40%	65%	90%*/88%**
Masimo PI Children	N.A.	35%	53%	41%*/31%**

Table 10.9: Summary of data from the PaedCEW study ($*$ denotes use of Newey-West 90% confidence interval with 200s lag, giving $p < 0.05$, $\star$ denotes exclusion of V7C, V13C, V14C, V16C, V8A, V12A from the dataset).

The MoorVMS laser Doppler device (Kalman filtered) clearly has the most consistent performance in the chamber cooling study, whilst the Masimo PI performs poorly in child volunteers, with erratic trend lines seen in the earlier plots. Improvements to the sensor fixation might improve the PI performance, for example the sensor could be secured firmly to the finger using adhesive tape. Although 80% of the adult volunteers saw a significant change in the PI (after obviously erroneous datasets were removed), it

can be seen from the earlier plots that the PI data are extremely noisy, and heavy low pass filtering is required to stabilise the baseline. This also means that, compared to the other techniques, the Masimo PI performs poorly in terms of baseline overlap and clear trend over the lowest 4°C of cooling.

Threshold based refill time recorded using the pneumatic capillary refill device was less consistent than the MoorVMS device, but arguably performed considerably better than Masimo PI. It may be notable that a trend over the lowest 4°C was seen more often than a non overlapping shift in the baseline refill time level. This observation, as well as the large "warm-up" drift in refill time at the start of the experiment (over one second in many recordings), suggests that the sensor may be influencing the refill time. Temperature recordings on a healthy volunteer demonstrate that the sensor is capable of cooling the skin surface be several degrees and warming by up to 1°C, confirming that this is a likely source of baseline drift. Standardisation of the room and sensor temperature during experiments, and a smaller sensor with reduced thermal mass may minimize these effects.

However, the 4 hour baseline recording from an adult volunteer (figure 10.49) shows short period refill time changes in the range of ±0.5s, apparently unrelated to temperature. These were removed from the 535nm refill time dataset by reducing the passband of the low pass filter from 2mHz to 0.5mHz (shown as the mauve dotted trace). This improved baseline stability at the expense of response time, and a clearer correlation with temperature is visible. The Masimo PI metric was low pass filtered using a 2mHz filter and the MoorVMS Kalman filter coefficients were selected so as to give an effective bandwidth of more than 2mHz with typical data (Kalman filter bandwidth is variable), so this change makes the capillary refill device considerably less responsive than either of the comparison devices, and only capable of outputting useful data approximately once every ten minutes. It should be noted that in typical clinical practice, refill time is measured at intervals of 15 minutes or longer, so this may still be a useful update rate.

10.6 Experimental characterisation of venous occlusion versus capillary refill

The protocol for this small scale experimental study was described earlier in section 10.1.1. The number of volunteers was much smaller than in the PaedCEW study, but, as described in detail earlier in this chapter; use of venous occlusion in conjunction with NIRS, and simultaneous application of a NIRS capable capillary refill device, together with a 5°C water bath to induce much more significant cooling of the forearm, allows

considerably more useful data to be gathered for the assessment of capillary refill as a perfusion metric.

The differential path approximation NIRS technique (as discussed in Chapter 8) was used to convert the five channels of raw intensity data into haemoglobin and oxy-haemoglobin molarity changes, leading to a (ideally minimal) fitting error at each wavelength. For plotting purposes the root mean square (RMS) fractional fitting error was calculated.

10.6.1 Proof of concept experiments on a healthy volunteer

Figure 10.55 shows a plot of 30 capillary refills of the Thenar Eminence of a healthy adult volunteer's hand. A pressure of 3PSI (155mmHg) was applied for a duration of four seconds every ten seconds. A forearm pressure cuff was used together with a modified Omron M2 base unit to apply 150mmHg pressure and occlude the arterial flow for ten of the capillary refills.

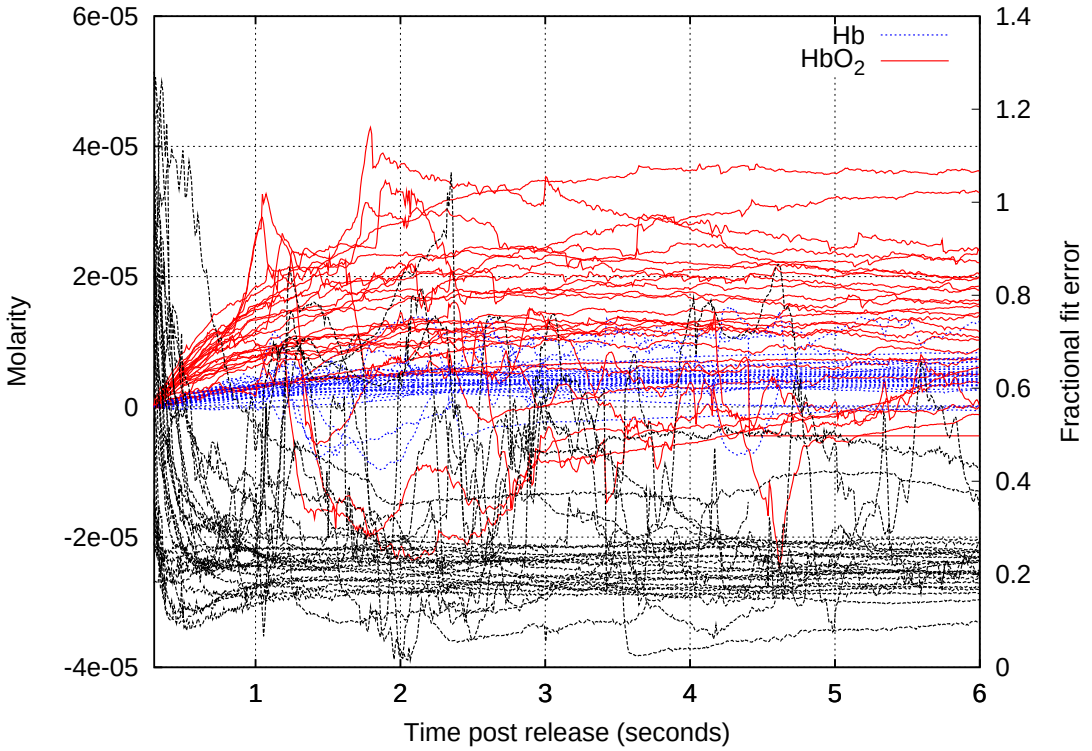


Figure 10.55: NIRS capillary refill data recorded from a healthy adult volunteer

The HbO_2 molarity traces (red solid), and all but one of the Hb traces (blue dashed) show an increase over the first seconds following release of pressure. The reference intensities for the differential path length spectroscopy technique were reset as the intensities at 0.3 seconds following pressure release during each refill.

The molarity rise during refill is split into two distinct clusters (although this is hard to make out in the figure). These correspond to the refills with and without application of arterial occlusion, with both the refill gradient and the total molarity change being lower when arterial occlusion was applied. However, the distinct clustering of traces starts to become less apparent further into the refill. Between two and six seconds post release, the molarity during occlusion appears to trend upwards, whilst the un-occluded data are stable. It is unclear from the figure if there is a significant change in the final molarity change between the occluded and non occluded refills, but answering this question may lead to an improved understanding of the capillary refill response, and is clearly worthy of further investigation.

Comparing the Hb and HbO_2 molarity changes in figure 10.55, it can be seen that the refilling blood has a saturation of slightly under 80%. This is consistent with published figures for typical venous oxygen saturation. The RMS fit error is around 0.2 for the majority of the refills, with some refills having a much higher and erratic fit error. Upon closer inspection of the raw intensity data, it was found that the sensor intermittently lost contact with the skin surface during some of the refills, causing large coupling errors and giving rise to significant fit error. The fit error can also be seen to be considerably higher and more erratic over the first 500ms or so following the release of blanching pressure. This may be due to the exaggerated effect of noise compared to the intensity change when the intensity change is small relative to the baseline level.

A line fit to the change in total haemoglobin molarity (i.e. $Hb+HbO_2$) over a short time range post release was then used as an approximate estimate of "effective capillary refill perfusion" by application of equation 10.1 derived earlier in this chapter to the gradient of a linear fit. The start point for the time range over which fitting was applied was chosen based on observation of the raw intensity data from both this recording and earlier results from Chapter 8, where it was found that there was an initial rapid and sometimes variable intensity shift over the first 0.1 to 0.3 seconds post release. The minimum time for return to baseline was always greater than 2 seconds, but after one second post refill the intensity profile often became more curved (this effect is evident in figure 10.55), so one second was selected as the endpoint. It was decided to use the averaged RMS fit error over this time range as a refill error metric to detect the degree to which loss of sensor contact affected the goodness of fit. The error estimates could then be used to calculate a weighted mean of these effective perfusion values.

With a weighted mean, equation 10.17 is no longer valid for estimation of the confidence interval. Assuming normalised weightings (i.e. the weighting coefficients w_n sum to 1):

$$\sigma^2 = \sum_{n=1}^{N} \left[w_n (x_n - \bar{x})^2 \right], \tag{10.27}$$

where $\bar{x}$ is the weighted mean of the set of effective capillary refill perfusion estimates. The variance of the mean can then be found as:

$$s^2 = \frac{\sigma^2}{1 - \sum_{n=1}^{N} w_n^2}, \tag{10.28}$$

with the 95% confidence interval then being the range $\pm 1.96s$ around the weighted mean $\bar{x}$. Equation 10.28 was used to produce a weighted mean together with 95% confidence interval for the two groups of refill gradients (with and without arterial occlusion).

Arterial occlusion	Weighted mean (with confidence interval)	Range
off	49.51 ± 29.68 ml/100g/min	16.2-25.6
on	19.80 ± 9.24 ml/100g/min	31.7-77.8
Difference	29.71 ± 31.08 ml/100g/min	6.1-61.6

Table 10.10: Summary of linear fit based perfusion derived from capillary refill data.

From the table it can be seen that the rate of molarity increase was consistently lower when arterial occlusion was used to halt the blood flow to the forearm. Although there is no overlap between the gradient ranges, the 95% confidence intervals calculated according to equation 10.28 overlap slightly. However, it should be noted from equation 10.19 that it is impossible to accurately estimate the confidence interval with such a small sample size.

Unfortunately it is difficult to apply arterial occlusion for a large number of consecutive capillary refills. The pressure cuff was found to be relatively uncomfortable at 150mmHg inflation pressure, and with a prolonged occlusion there is a risk of tissue damage. For these reasons it was judged to be unsafe to attempt a long sequence of capillary refills with arterial occlusion.

Although the data do not demonstrate a statistically significant change in the rate of tissue total Hb molarity increase during capillary refill, the 95% confidence interval in table 10.10 is only slightly larger than the difference in mean gradients. For this reason it would seem reasonable to conclude that the arterial occlusion had the effect of reducing the rate at which blood refilled the capillary bed following the pressure release. This would support the conclusions reached from analysis of the PaedCEW data; that there is a connection between perfusion and the capillary refill response.

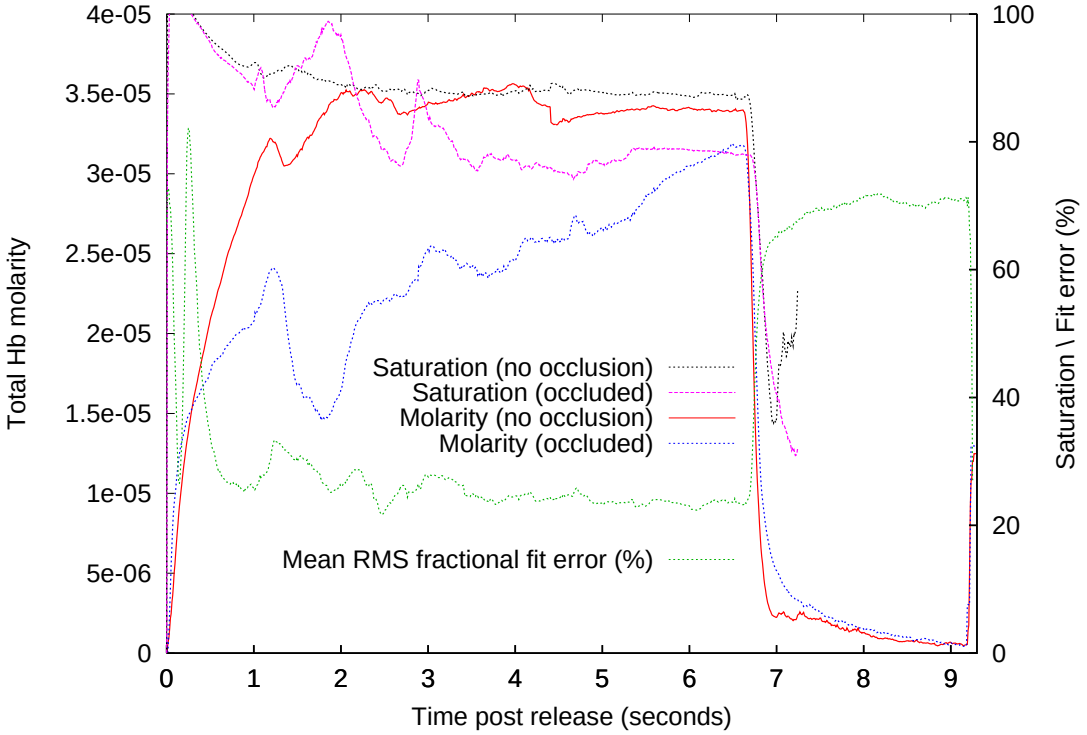


Figure 10.56: NIRS derived total haemoglobin (Hb) molarity and saturation versus time with and without simultaneous arterial occlusion. Saturation was only plotted to 7.5 seconds for clarity due to heavy noise after this point.

The mean total haemoglobin molarity was calculated as a function of time from the volunteer dataset (shown in figure 10.55), in an attempt to further understand the differences in NIRS capillary refill data with and without arterial occlusion. This is shown as the two total molarity traces in figure 10.56. For clarity in this plot, the reference intensities were set as the intensities at the end of the pressure application before each refill. It can be seen that there are apparently erroneous fluctuations in the molarity traces (e.g. at 1.5 seconds post pressure release). These are the same erratic fluctuations which were earlier found to be due to variable sensor to skin contact from inspection of the raw intensity data. The amplitude of the artefacts is greatest in the occlusion trace. This is probably due to the volunteer recording containing only ten refills with the arterial occlusion applied, compared to twenty un-occluded refills (ten before occlusion and ten following).

From the plot it can be seen that the overall pressure induced change in total molarity is approximately the same with and without occlusion of the arterial flow, this was not clearly discernible from figure 10.55. It would appear that the molarity during occlusion is still rising at six seconds post release, and so would be likely to reach or potentially overshoot the non occluded molarity level if the time between pressure applications was

increased, but this would come at the cost of fewer refills being possible in the limited occlusion time, so was not attempted. It may also be notable that the two traces are very similar over the first 0.3 seconds. This is an intriguing finding, as the time period approximately 0.3 seconds post release was identified in Chapter 8 as being an apparent change point in the intensity curve; the point where an initial rapid decrease in intensity, and one that had significant inter site variability, gave way to a slower and more consistent semi-linear or curved refill response in a number of the early volunteer recordings. It should be noted that the pressure application device (development of which was discussed in Chapter 8) releases pressure very rapidly, taking between 50 and 100 milliseconds to return to a baseline pressure of 0.2PSI. The mean RMS fit error trace would appear to confirm the hypothesis that some form of structural or geometric change takes place in the tissue; after the pressure is released, the error initially falls, before peaking again at approximately 0.25 seconds post release. The fit error decreases to 23% as the blood refills the capillary bed between 0.3 and 6 seconds post release, suggesting that the differential path length approximation is a better approximation to the refill than to the pressure application. Finally, the fractional fit error is very high during the pressure application period. This would be expected, as the intensities are very close to the reference intensities during this period, meaning small errors (for example due to coupling) have a large effect.

The separation between the two traces suggests that multiple mechanisms are responsible for the NIRS intensity change during capillary refill. The saturation data support this conclusion (this was defined as the oxy-haemoglobin molarity as a fraction of total haemoglobin), as the refill is fully saturated (i.e. the NIRS processing results in an increase in oxy-haemoglobin only) over the first 0.25 seconds, with no difference in saturation between the occluded and non-occluded refills prior to 0.3 seconds. The final saturation at six seconds post release is approximately 78% during occlusion and 85% for the non-occluded refills. Saturation would be expected to fall during occlusion as there is no supply of fresh oxygen to the tissue.

10.6.2 Discussion of the NIRS data during capillary refill

The features seen over the first 0.3 seconds in figure 10.56 are hard to explain in terms of refilling venuoles or arterioles. Intuitively it might be expected that these vessels would refill rapidly once the blanching pressure was released, but in the case of refilling venuoles the initial saturation would be very similar to the final saturation, and in the case of refilling arterioles, the saturation might be close to 100% without occlusion, but would be expected to be similar to the final saturation of ∼78% during occlusion as

the static blood pooled within the arm desaturated. An alternative explanation for the first 0.3 seconds may be that pressure application changes the optical properties of the tissue, for example via movement of interstitial fluid. However, in simple models this effect would simply lead to an increase in μ_s' whilst blanching pressure was applied, and thus to a reduction in μ_{eff} as pressure was released. More complex effects involving both distortion of the layered structure of the skin and fluid transport may still be a feasible explanation for the NIRS traces.

In clinical studies the capillary refill is often explained in terms of refilling of the capillary bed directly by perfusing arterial blood. This explanation cannot account for the very high rate of tissue molarity increase, which was earlier found to correspond to a mean perfusion of approximately 50ml/100g/minute in the healthy volunteer. This is over an order of magnitude higher than both typical values in published studies, and the value derived earlier from NIRS and venous occlusion applied to the same volunteer.

An alternative theory of capillary refill is that the blood refills from the venous side, but this explanation fails to account for the reduction in refill rate induced by arterial occlusion. There are also flap valves within the veins which prevent backflow through the larger subcutaneous vessels.

One theory which does match the refill response is that interconnected capillary vessels within the capillary bed itself are a significant source of refilling blood. In the case of a perfused capillary bed in the region surrounding the blanched area, the mean blood pressure within this region would be close to its normal level (there may be some disturbance due to the 0.2PSI/10mmHg baseline sensor application pressure), whilst with occlusion applied this pressure would be closer to zero, leading to a reduction in the refilling blood flow into the blanched region from surrounding capillary bed.

In order to ascertain if this is a feasible explanation, it is necessary to consider both the degree of interconnectivity within the different classes of blood vessels, and the typical mean pressures within them.

Published studies of capillary bed pressure have found values in the range 10.5mmHg to 55mmHg. Shore used a non-invasive capillaroscopy technique [264], finding nailbed capillary pressures in the 10.5 to 22.5mmHg range in a group of healthy adult volunteers. Levick et al used an invasive method; a scalpel was employed to remove a small area of skin from a finger or toe, allowing a micro-capillary to be directly inserted into the capillary bed to measure the internal blood pressure [265]. Arterial and venous pressure was recorded from the volunteers in the Levick study, allowing the pressure drop through the micro-vascular system to be accurately quantified. It can be seen (figure 4 in the publication) that the elevation relative to the heart has a large influence

upon pressures in all vessels, as might be expected. However, as the volunteer's arm was positioned horizontally on the surface of a desk at the same elevation as the heart in the capillary refill experiment, it would appear that this effect was negligible. A similar protocol was followed in the PaedCEW study to avoid biassing of the results due to elevation variations. A second significant observation from the figure is that venous pressures were centred around 6mmHg when the hand was at equal elevation to the heart. This pressure may be sufficient to cause backfilling of the capillary bed, but it is lower than the 10mmHg baseline application pressure of the capillary refill device, so the venous compartment would not appear to be a significant blood source during refill. The capillary pressures recorded in the hand by Levick were higher than in the study by Shore, with the mean capillary pressure of 33mmHg. The reasons for this discrepancy are unclear, but may be related to the non-invasive method of pressure measurement by Shore.

Mean arterial pressure of the group of healthy adult volunteers in the Levick study was 80mmHg, so compared to this, 33mmHg mean capillary bed pressure is sufficiently high to cause significant capillary blood flow during the refill, provided that the vascular resistance to flow between neighbouring capillary beds is lower than that between arteries and veins.

The existence of significant interconnectivity within the dermal capillary bed has been established by numerous authors, Braverman et al give an overview of developments in theoretical understanding of the micro-circulatory vessel topology [266]. However, it is not immediately clear if the degree of interconnectivity within the capillary bed is sufficient to explain the very high effective perfusion during capillary refill, or how the interconnectivity between neighbouring capillary beds differs with their physical separation for each class of blood vessel.

It is unclear if sufficient data on interconnectivity within the human dermal capillary bed exists for a definitive answer to be reached. For example, in his study of scaling laws in vascular trees [267], Kassab states that "the use of a symmetric tree model is necessitated by the fact that data on the realistic connectivity or asymmetry of most vascular trees considered are lacking". Significant interconnectivity is not always a feature of the micro-circulation, for example, in a study of Tibialis anterior muscles in the hamster [268], Lund et al found capillaries organised into "modules"; small subunits of less than a cubic mm in volume with no interconnectivity to surrounding capillaries. Berg describes such capillary organisation as a general feature of striated muscle [269].

However, muscle tissue (and striated muscle in particular) is considerably more ordered than the skin and subcutaneous tissue, so it would appear that the capillary structures

described by Berg and Lund lie close to the extreme case of repeating capillary subunits.

The hypothesis that refilling blood flows through the capillary bed could be tested by measuring the change in refill induced by skin cooling. As this will lead to vasoconstriction, the rate of molarity increase during refill might decrease both with and without arterial occlusion if flow from the neighbouring vascular bed is the principle source of refilling blood, but only if the flow involves sufficiently high order vessels that smooth muscle tissue or precapillary sphincters have a significant influence.

10.6.3 Experimental characterisation of capillary refill versus venous and arterial occlusion with 10°C water bath cooling

The protocol for this experiment was discussed earlier in subsection 10.1.1. It was repeated with a number of healthy adults volunteers recruited within the university. All volunteer recordings (both in this subsection and in earlier subsections) were subject to the university ethical review process, and written consent was obtained from all volunteers. It was found that it was impossible to apply consistent cuff pressures with manual operation of the Omron M2 device. For this reason a control board from one of the data loggers (Chapter 9) was used to drive pneumatics from an Omron pressure cuff in these experiments. Unfortunately the pressure control loop in the data logger firmware was not well suited to the Omron hardware, resulting in a slow pressure ramp up at the start of each occlusion (∼25 second ramp up as opposed to ∼10). This may well have increased venous pressure, biassing the capillary refill data.

Healthy adult volunteer 1

This experiment had 16 minutes total duration, 5 minutes 30 seconds of which was spent with the subjects forearm placed in the water bath. Water temperature was measured as 11.5°C prior to the experiment, and reached 12°C at 16 minutes.

Key features seen in the volunteer PI plots are annotated in figure 10.57; the vertical dashed line at 6 minutes indicates the point at which the forearm was placed into the water bath. A clear decrease in PI is seen over the following 3 minutes of water bath exposure, with the PI then maintaining a level around 1 for the remainder of the water bath exposure and for several minutes afterwards. The arm was withdrawn from the bath at 11.5 minutes, visible as the large upward spike in PI, likely due to motion artefacts. A similar motion effect occurred as the arm was placed into the bath, here the pulse oximeter reported a loss of signal, indicated as a discontinuity in the trace.

The two periods when arterial occlusion was applied are clearly visible at 2minutes and 13minutes into the experiment. Occlusion was applied using a blood pressure cuff above

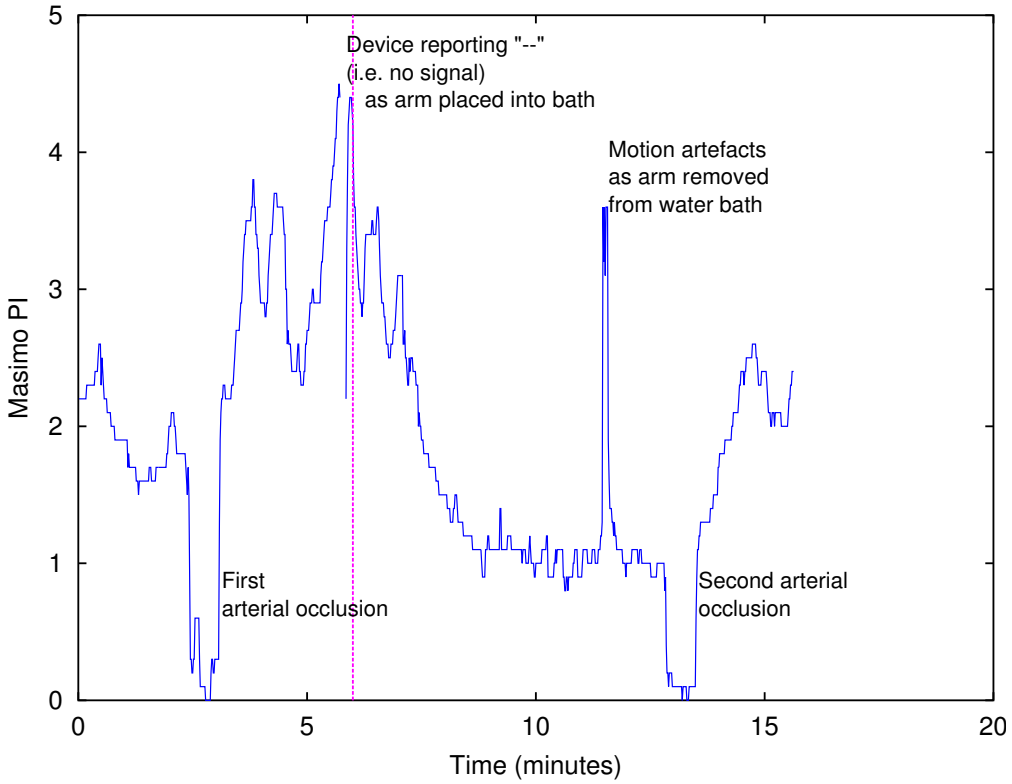


Figure 10.57: Masimo PI recorded from finger of the first healthy adult volunteer.

the elbow inflated to 145mmHg for 1 minute. A large drop in PI occurred as pressure was applied; PI fell to zero in both occlusions with 90% of the fall occurring in less than 10 seconds, suggesting that any low pass filtering carried out by the Radical 7 unit has a cutoff frequency higher than 0.1Hz, at least once the device has been operating for some minutes. The experiments in section 10.4.2 suggest that adaptive filtering is used.

Figures 10.58 and 10.59 show the NIRS molarity during capillary refills. As the Omron M2 cuff controller took approximately 20 seconds to reach the 150mmHg pressure required for arterial occlusion, two capillary refills recorded in this period were excluded.

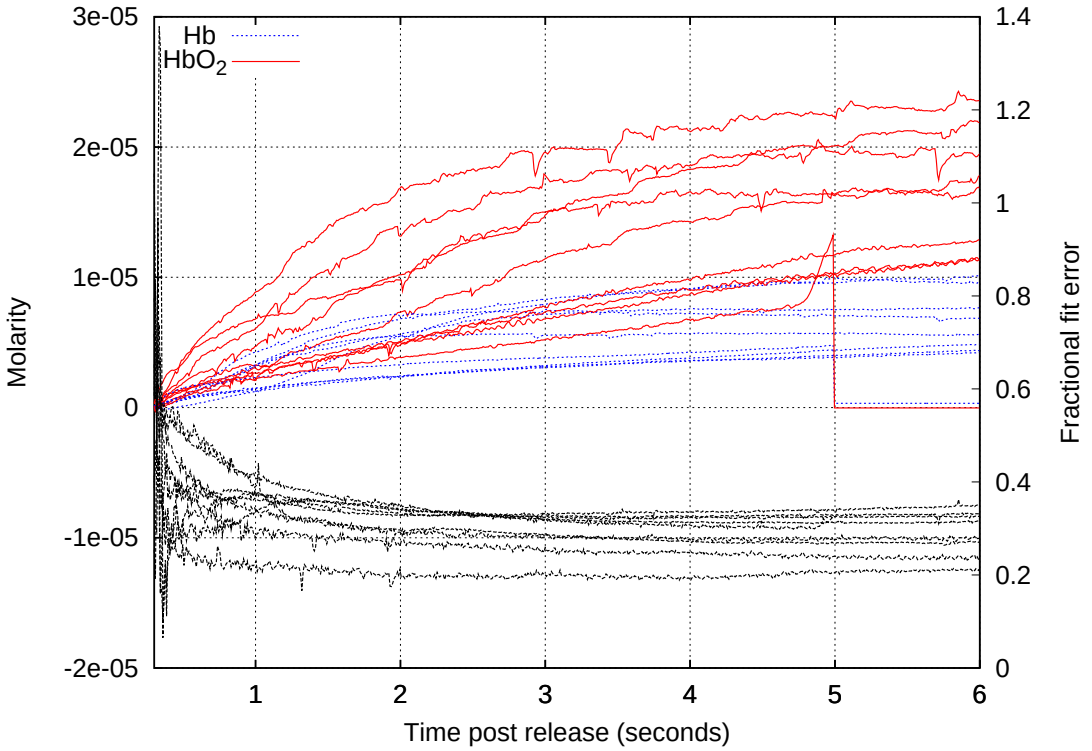


Figure 10.58: Raw NIRS based molarity recorded during room temperature capillary refills, with and without arterial occlusion.

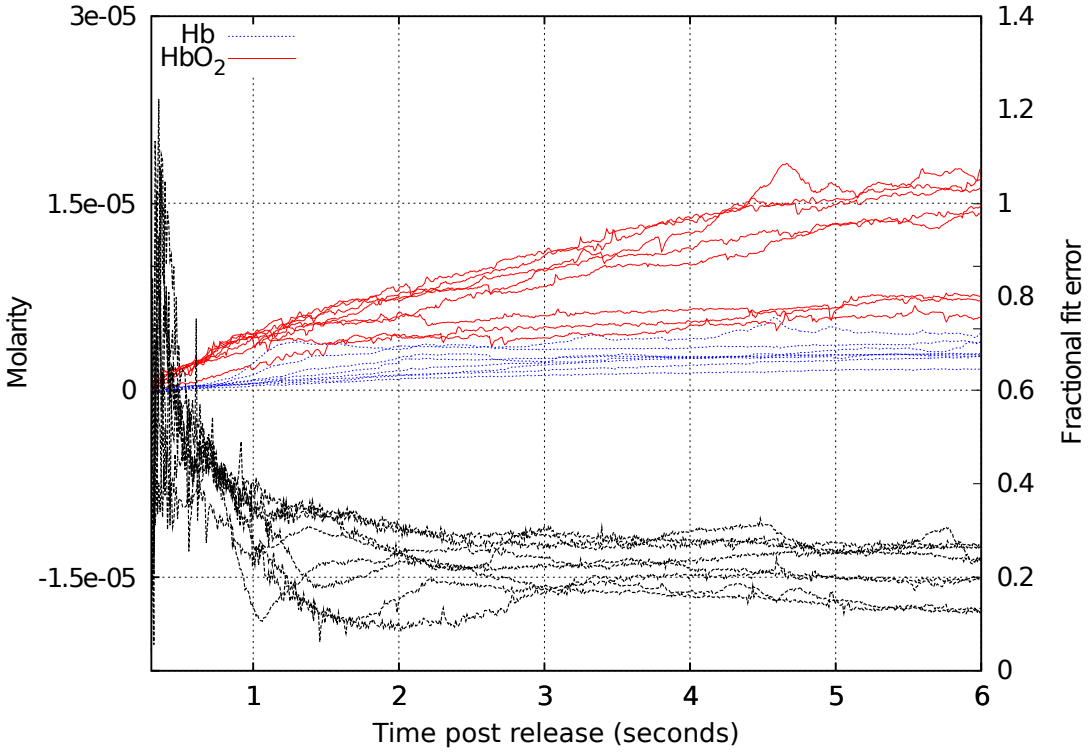


Figure 10.59: As above, recorded during post water bath capillary refills.

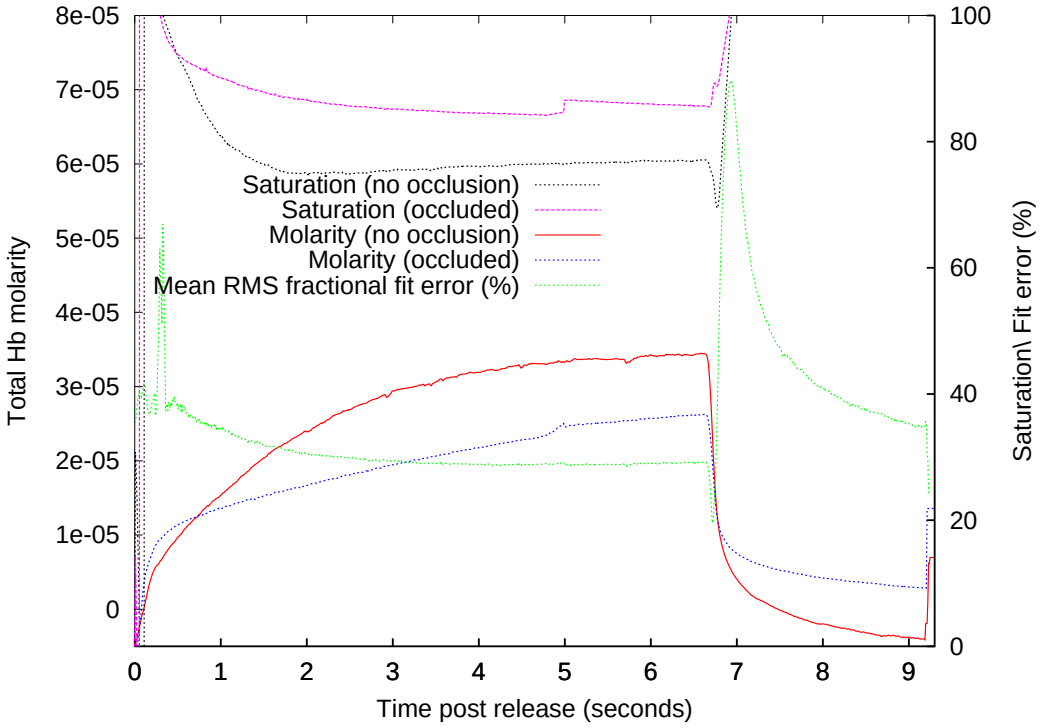


Figure 10.60: Weighted mean NIRS based molarity, error and saturation recorded during room temperature capillary refills, with and without arterial occlusion.

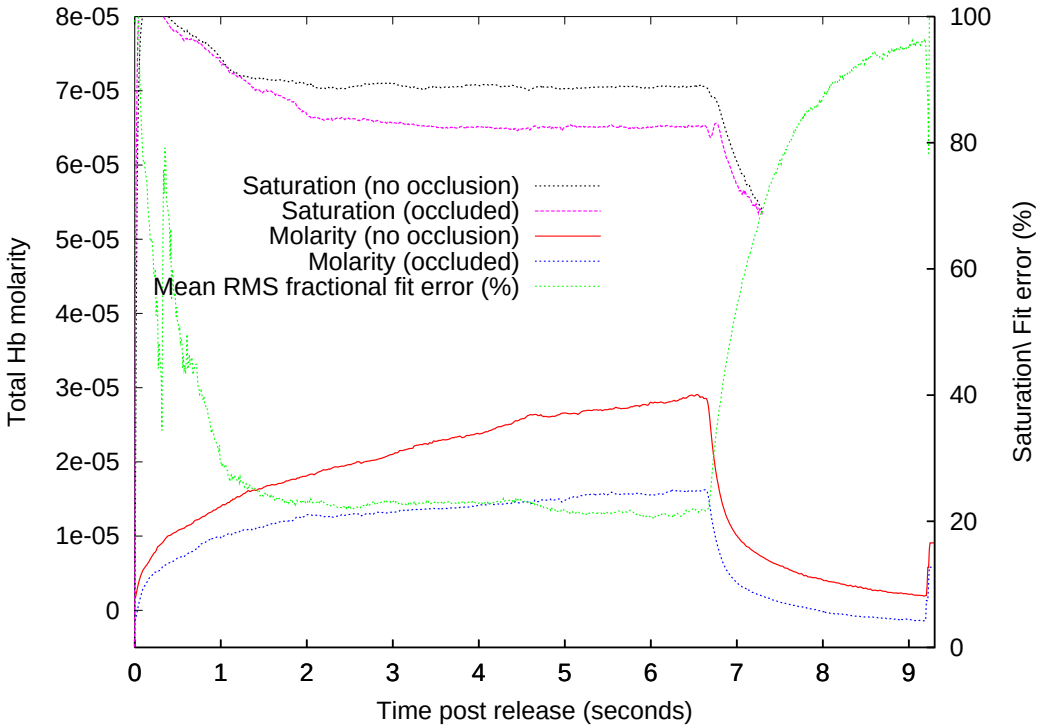


Figure 10.61: As above, recorded during post water bath capillary refills.

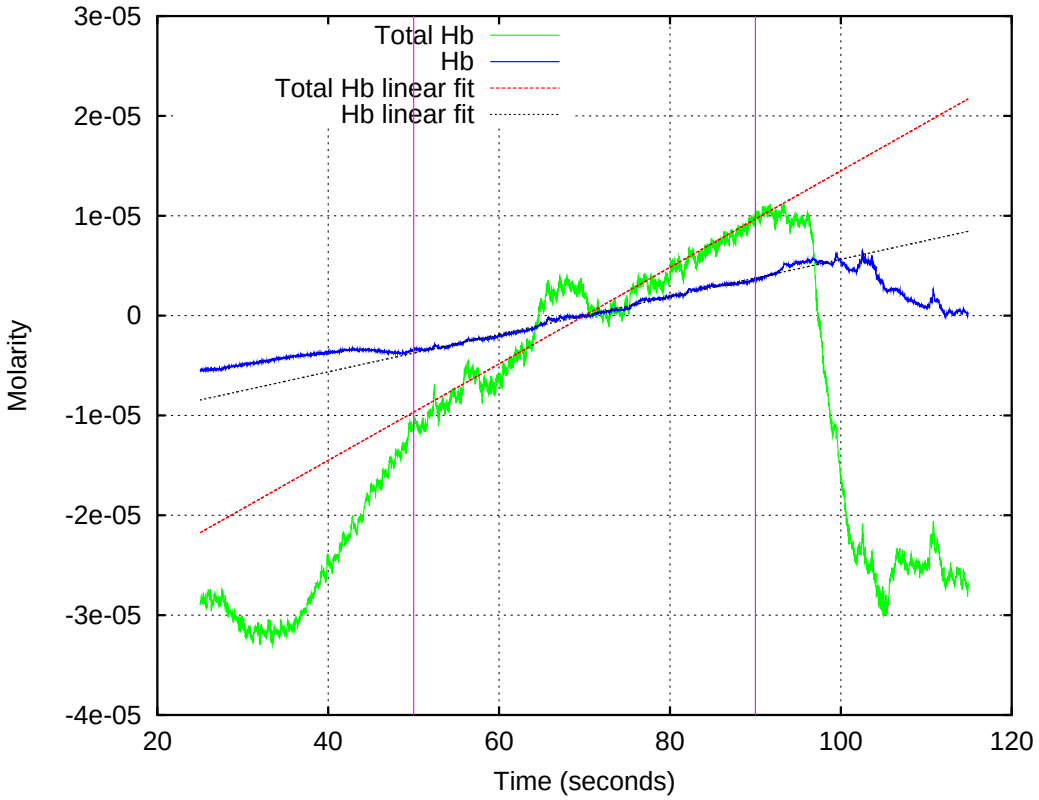


Figure 10.62: Venous occlusion NIRS recorded at room temperature.

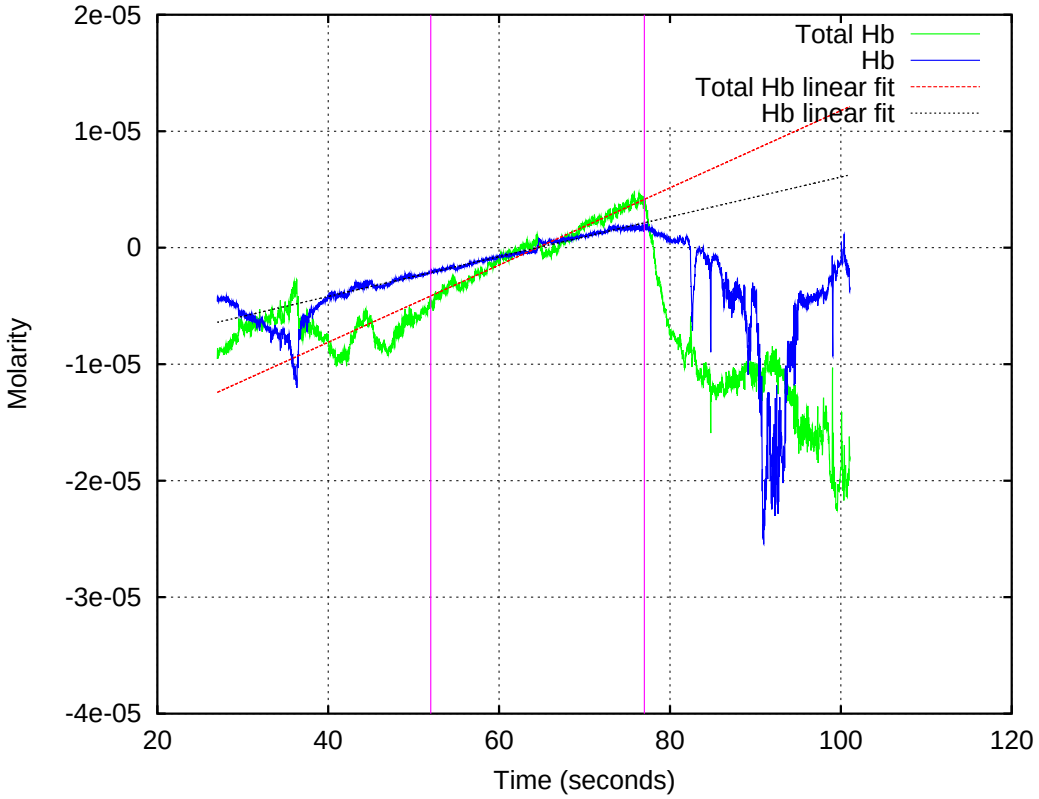


Figure 10.63: As above, post water bath.

Experiment	Perfusion (ml/100g/min)	Venous saturation (%)	Metabolic rate (mlO_2/100g/min)
Room temperature	1.26 ±0.14	61.1 ±5.5	0.115 ±0.010
Water bath cooled	0.86 ±0.11	48.5 ±7.8	0.104 ±0.006

Table 10.11: Fit results and errors from the venous occlusion/NIRS.

Table 10.11 presents the NIRS results of the pre and post water bath venous occlusions, the errors shown are 95% bounds derived using the Newey-West method described in section 10.2.1. The raw NIRS molarity data can be seen in figures 10.62 and 10.63. There was a statistically significant decrease in volumetric perfusion between the room temperature and post water bath states, but no significant change in the metabolic rate of oxygen consumption (this would be expected to be constant). The rate of oxygen consumption was less than half that recorded from an different volunteer in earlier work (Table 10.3), but was within the range of the values found by Zurlo et al [249], unlike the earlier experiment. The Masimo PI data indicate that the cooling induced perfusion reduction effect began to diminish rapidly after approximately 13 minutes into the experiment, meaning that the venous occlusion measurement is likely to have underestimated the reduction in perfusion.

Considerable difficulties were encountered with skin contact; large artefacts caused by varying sensor skin contact can be seen in figure 10.63 after the occlusion period. A improved pressure control system may help here.

Figures 10.60 and 10.61 show processed NIRS data from the first volunteer following a weighted averaging process (described in detail in section 10.6.1). It can be seen that cooling resulted in a reduction in both the total molarity change during refill, and the rate of total molarity change immediately post release when no arterial occlusion was applied. The effect of arterial occlusion was similar to that seen earlier in a different volunteer (section 10.6.1); the rate of change of total molarity following pressure release was reduced considerably.

There was only a slight cooling induced difference between the two sets of refills conducted with arterial occlusion; the total change in molarity decreased slightly with cooling, but there was no clear difference in the rate of total molarity increase over the 0.3 to 1 second period post release. Over the 1 to 7 second period post release, the room temperature molarity appears to follow an approximately linear trend with time, whereas the molarity in the cooled hand stabilised at 15μM. However, due to the small number of refills (4 room temperature and 3 cooled refills with arterial occlusion), only limited conclusions can be drawn from the data.

Healthy adult volunteer 2

This healthy adult volunteer experiment had a total duration of 18minutes, 5 of which were spent with the subject's forearm immersed in the water bath, which was at 9°C at the start of the immersion and 10°C after the experiment. The subject's forearm was placed into the water bath 7 minutes into the experiment (indicated with the dashed vertical magenta line).

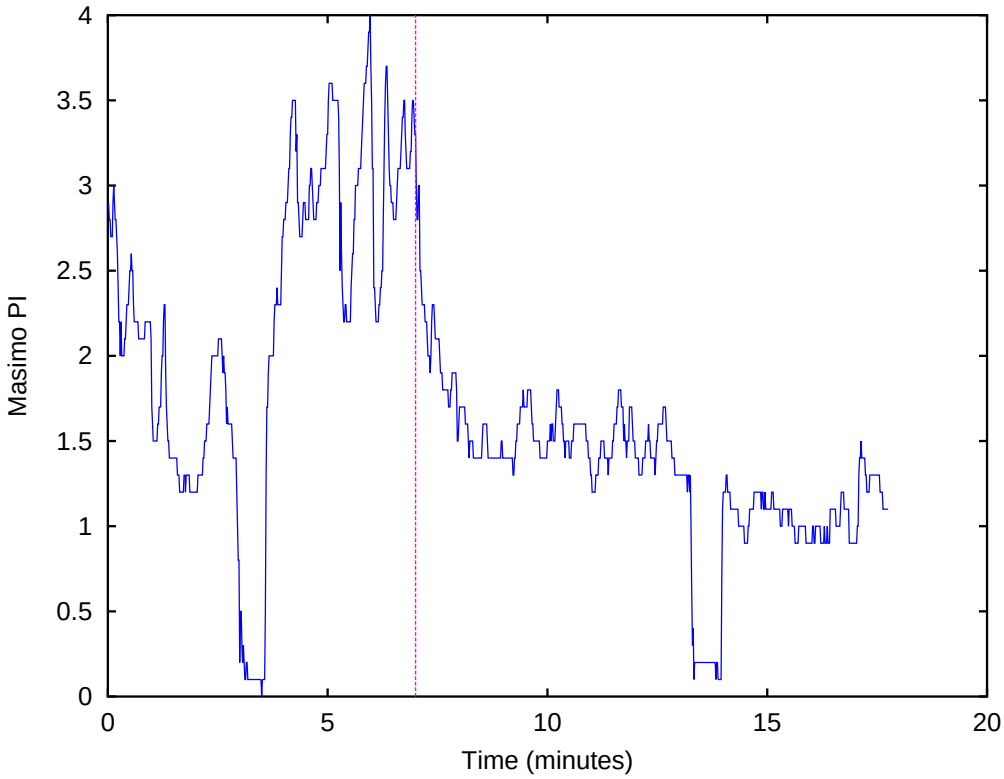


Figure 10.64: Masimo PI recorded from finger of the second healthy adult volunteer.

From figure 10.64, it can be seen that the PI is unstable prior to cooling, varying between 1.2 and 4. There is a pronounced drop in PI over the first minute of cooling, although the cooled PI baseline (∼1.4) is slightly higher than the lowest PI value reached prior to cooling. The PI continued to fall after the forearm was taken out of the water bath, suggesting that the cooling response may have been delayed in the second volunteer.

NIRS molarities during the capillary refills are shown in figures 10.65 to 10.68, and raw venous occlusion data in figures 10.69 and 10.70. The processing used to produce all figures was the same as that applied to the previous dataset.

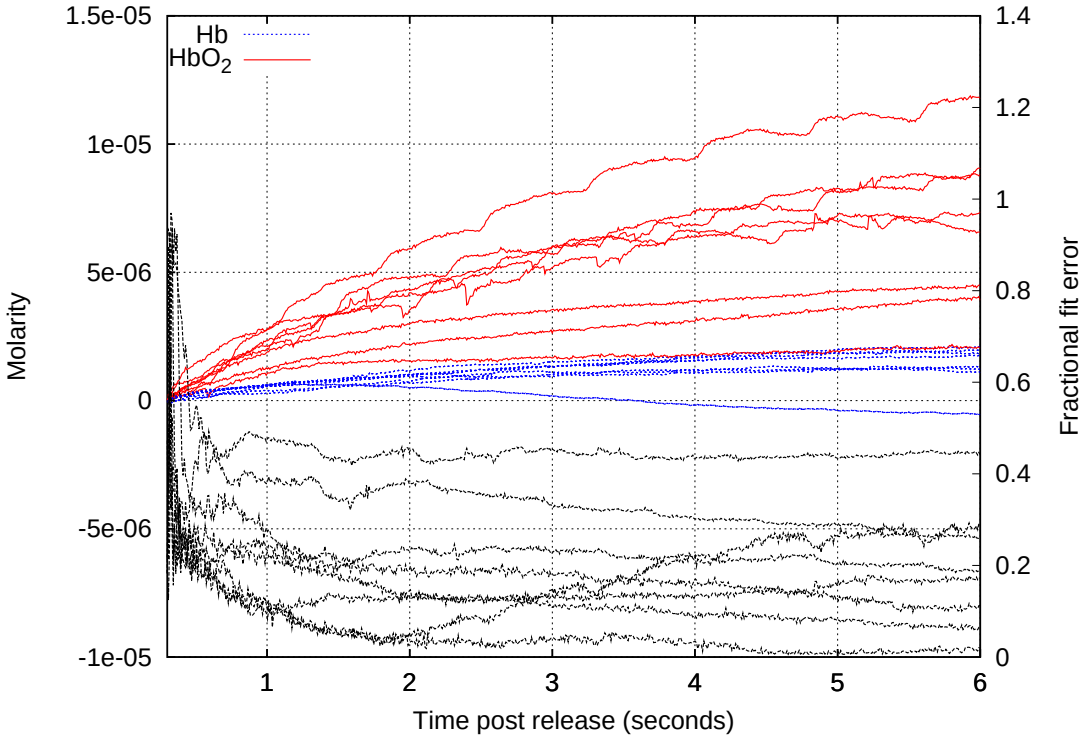


Figure 10.65: Raw NIRS based molarity recorded during room temperature capillary refills, with and without arterial occlusion.

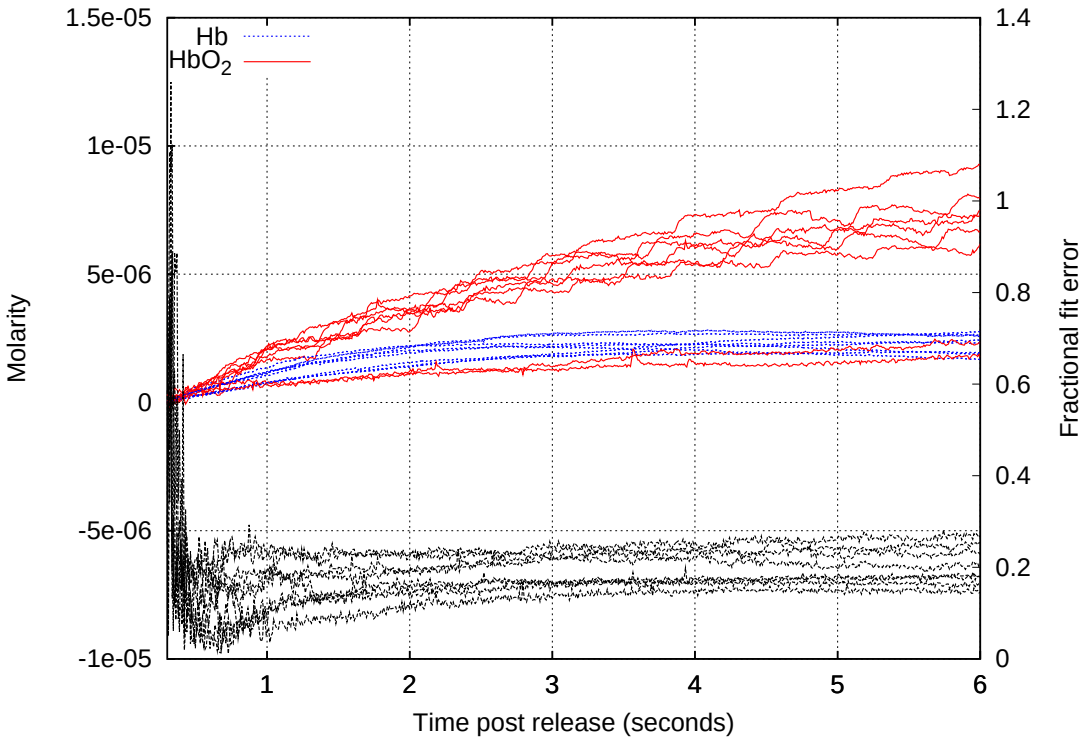


Figure 10.66: As above, recorded during post water bath capillary refills.

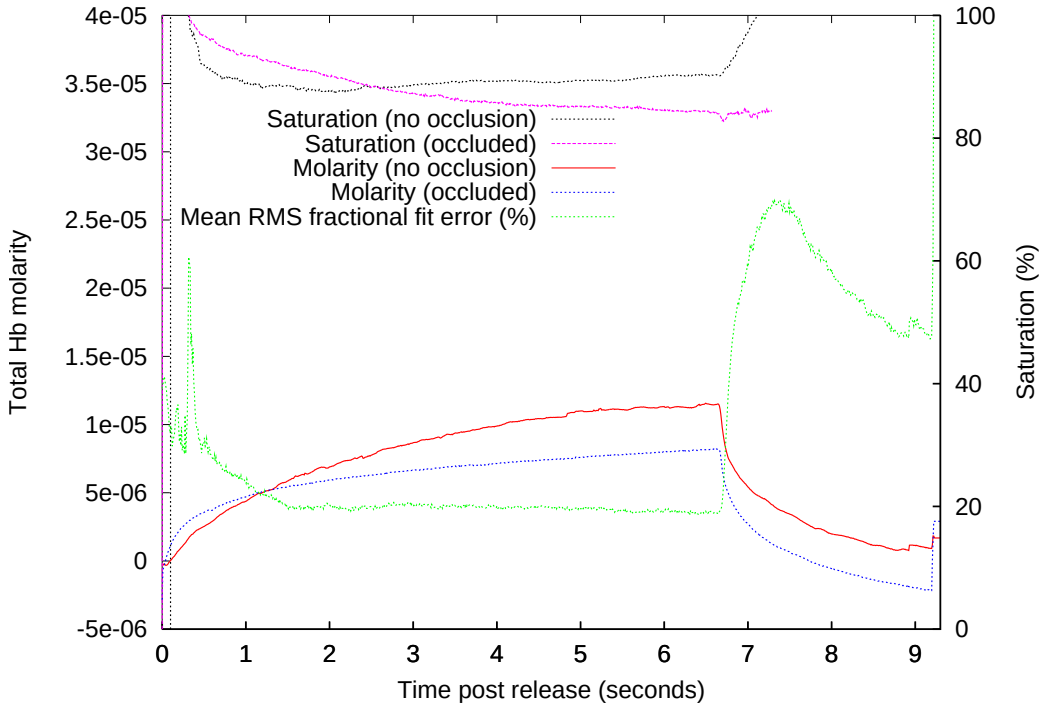


Figure 10.67: Weighted mean NIRS based molarity, error and saturation recorded during room temperature capillary refills, with and without arterial occlusion.

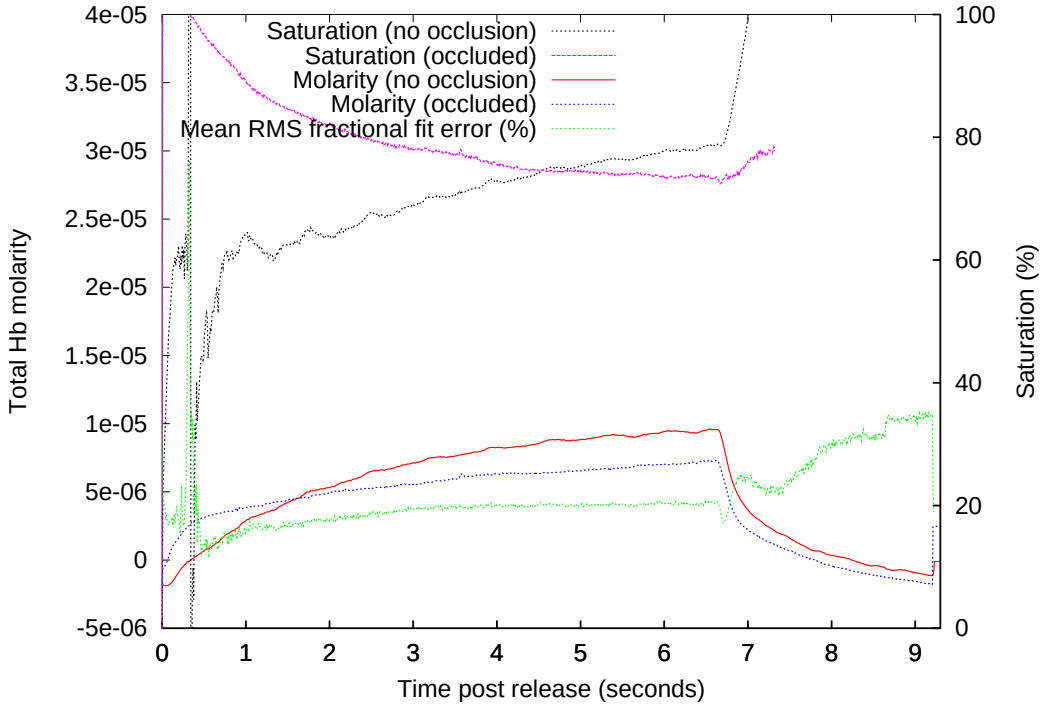


Figure 10.68: As above, recorded during post water bath capillary refills.

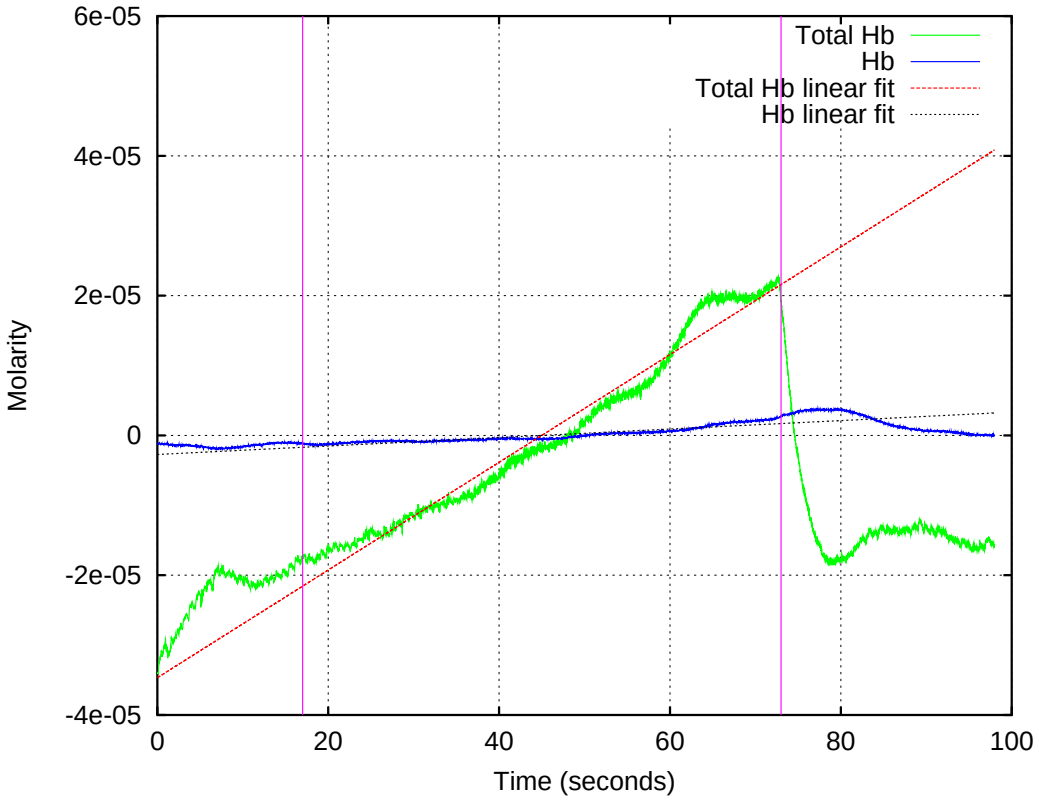


Figure 10.69: Venous occlusion NIRS recorded at room temperature.

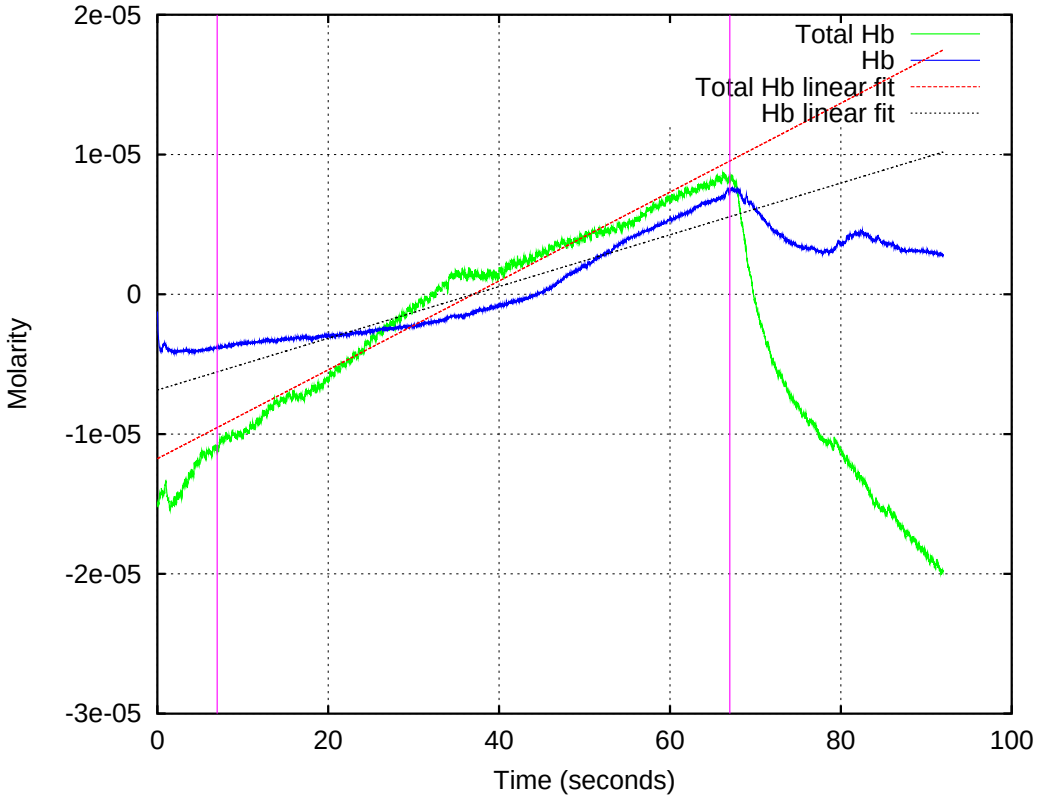


Figure 10.70: As above, post water bath.

Experiment	Perfusion (ml/100g/min)	Venous saturation (%)	Metabolic rate (mlO_2/100g/min)
Room temperature	2.01 ±0.30	92.1 ±5.1	0.037 ±0.021
Water bath cooled	0.83 ±0.17	41.8 ±17.2	0.113 ±0.023

Table 10.12: Fit results and errors from the venous occlusion/NIRS.

Table 10.12 presents the volumetric perfusion, venous saturation and tissue metabolic rates of oxygen derived via venous occlusion. Following forearm cooling there was a very significant drop in perfusion, but also a significant increase in metabolic rate, which would not be expected. Closer analysis of figure 10.69 suggests that the room temperature metabolic rate calculation was biassed by the relatively high perfusion seen in this volunteer; as the rate of HbO_2 molarity increase was much higher than that of Hb, a small error in the path length model may easily have lead to a large enough negative Hb bias from HbO_2 to explain the result. The fact that the metabolic rate in the cooled state (0.113 mlO_2/100g/min) is close to published values supports this hypothesis.

Healthy adult volunteer 3

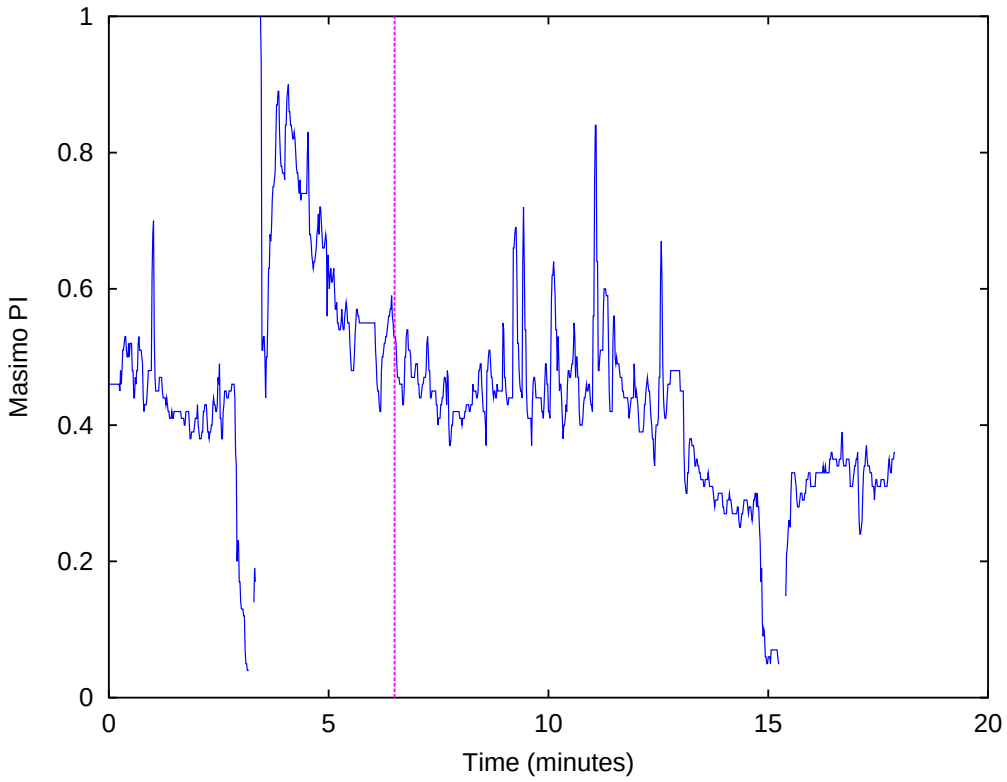


Figure 10.71: Masimo PI recorded from finger of the third healthy adult volunteer.

Figure 10.71 shows similar PI trend behaviour to the previous volunteer. There are "signal drop outs" during both arterial occlusions, visible as discontinuities in the trace. Water bath temperature was 10°C, increasing to 11°C by the end of the recording.

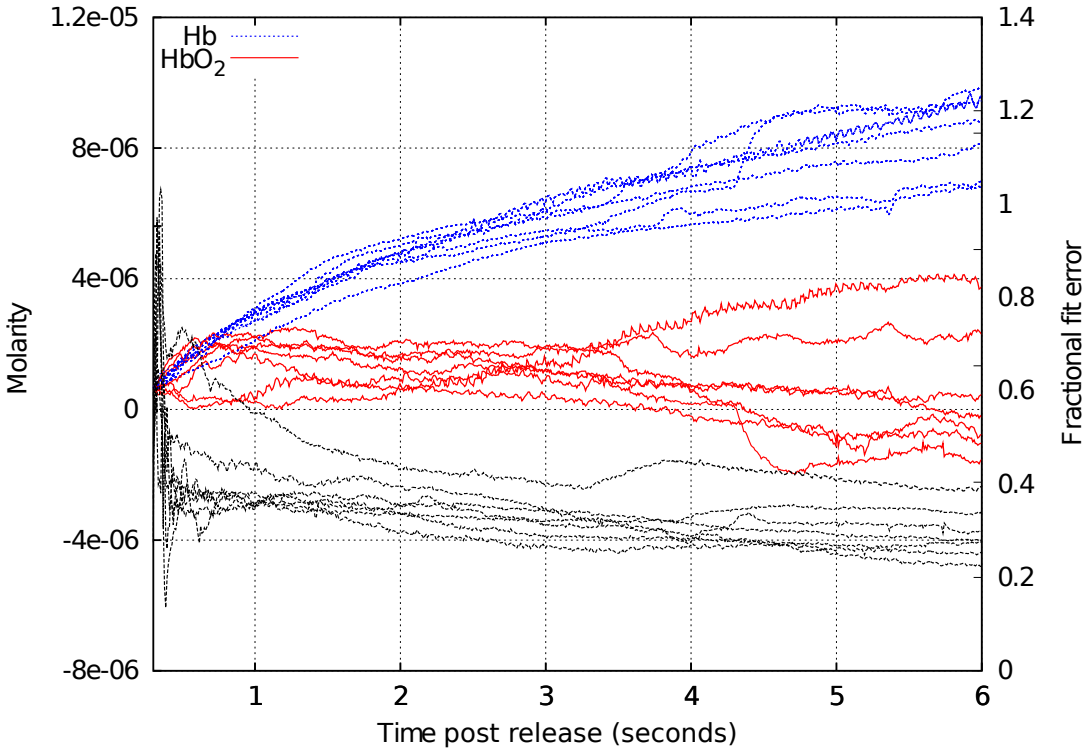


Figure 10.72: Raw NIRS based molarity recorded during room temperature capillary refills, with and without arterial occlusion.

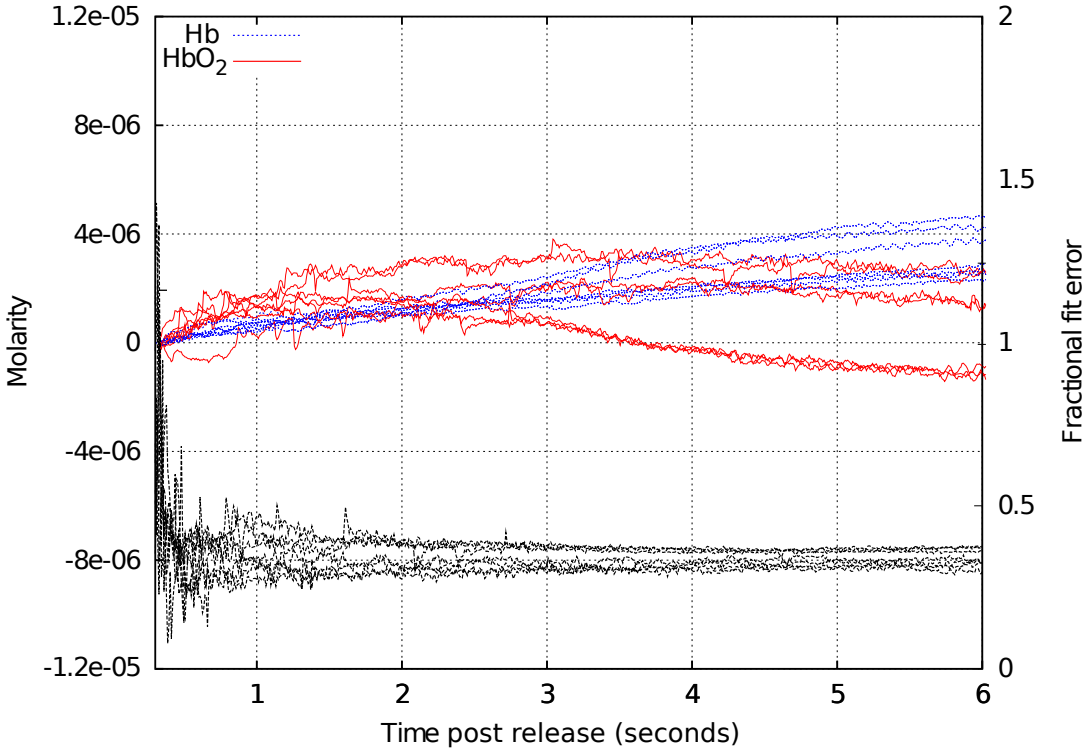


Figure 10.73: As above, recorded during post water bath capillary refills.

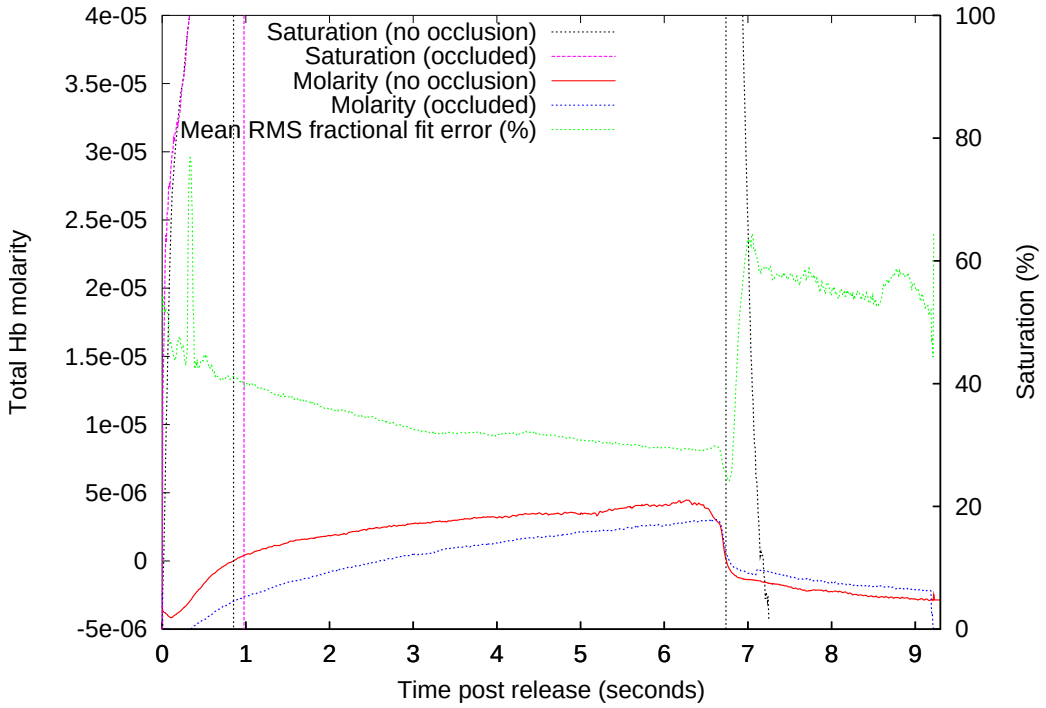


Figure 10.74: Weighted mean NIRS based molarity, error and saturation recorded during room temperature capillary refills, with and without arterial occlusion.

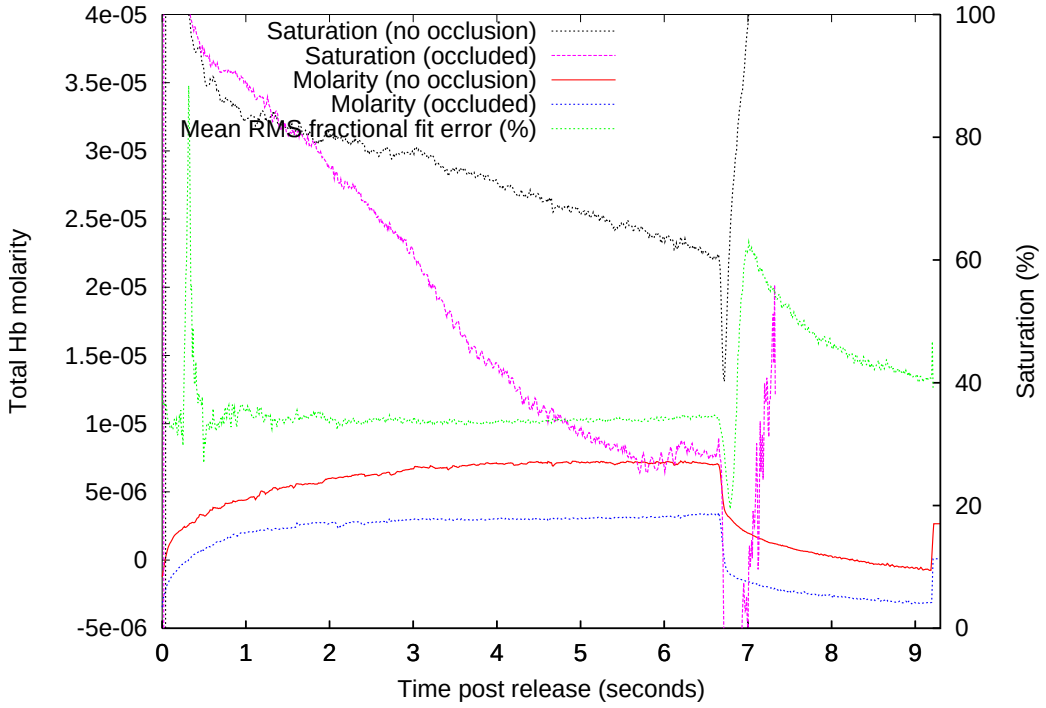


Figure 10.75: As above, recorded during post water bath capillary refills.

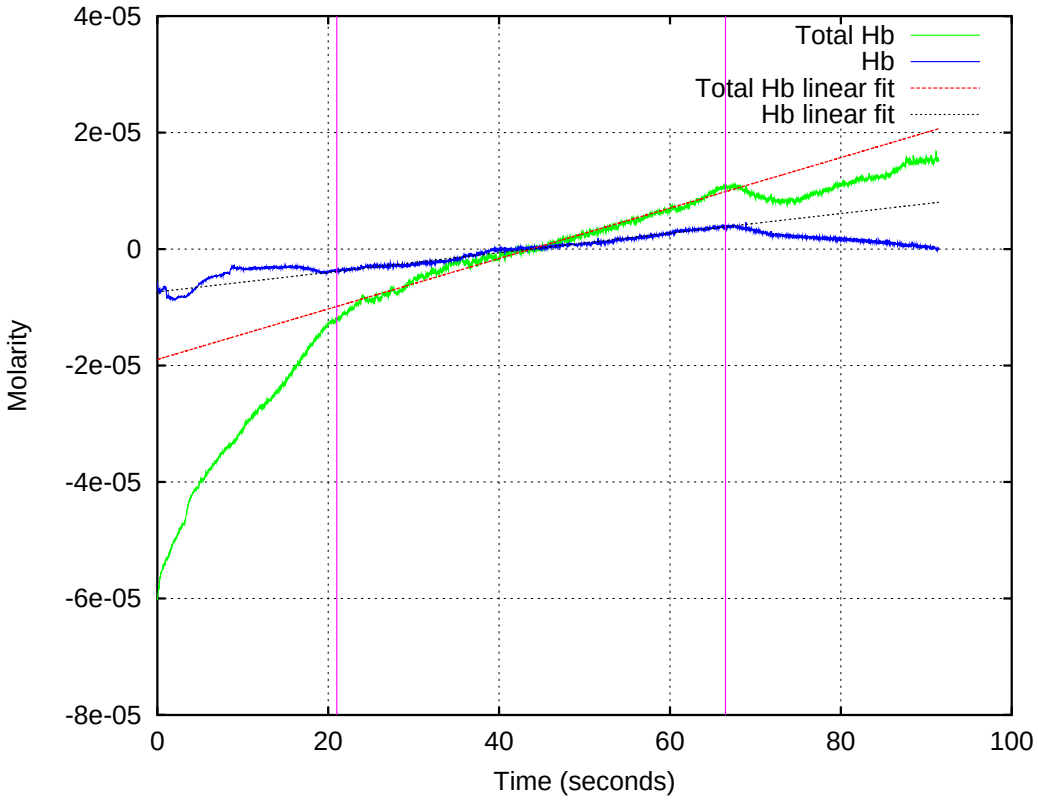


Figure 10.76: Venous occlusion NIRS recorded at room temperature.

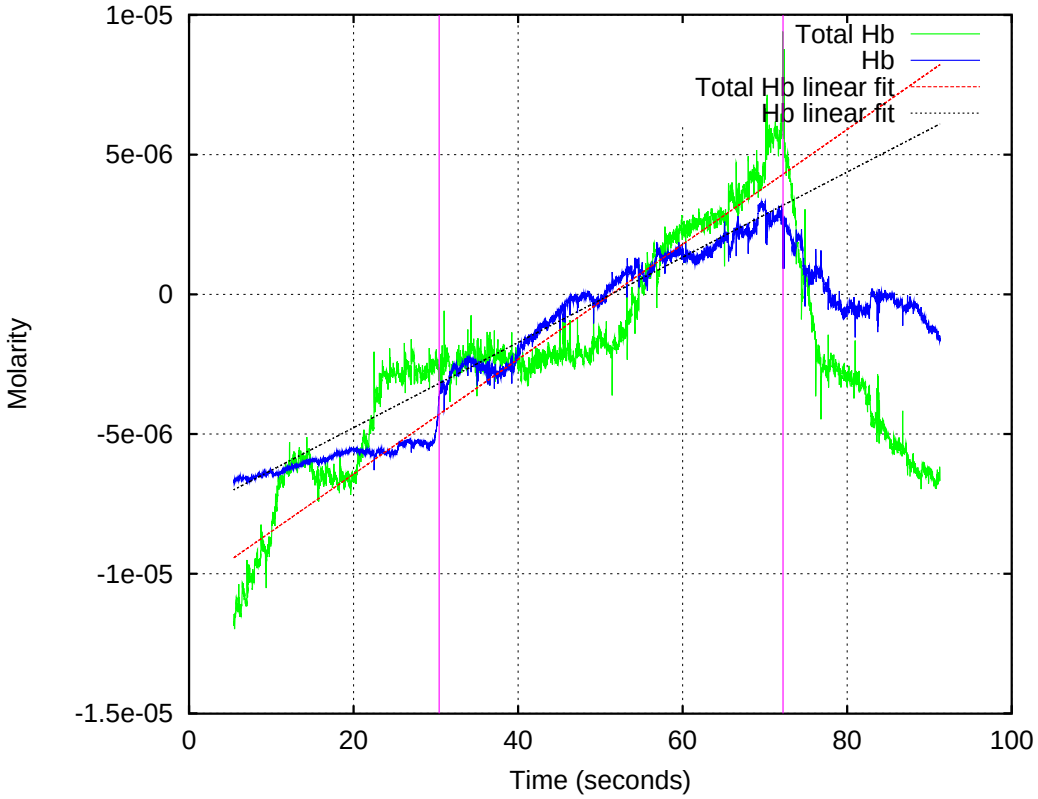


Figure 10.77: As above, post water bath.

Experiment	Perfusion (ml/100g/min)	Venous saturation (%)	Metabolic rate (mlO_2/100g/min)
Room temperature	1.13 ±0.31	61.2 ±14.9	0.103 ±0.014
Water bath cooled	0.54 ±0.37	25.8 ±60.5	0.093 ±0.015

Table 10.13: Fit results and errors from the venous occlusion/NIRS.

Good venous occlusion data were obtained from this volunteer, with a significant decrease in volumetric perfusion, no significant change in metabolic rate, and a similar metabolic rate to both the previous volunteers and published data. However, the capillary refill NIRS data had a high RMS error, and indicated an excessively low, or even negative oxygen saturation. The reason for the poor performance of NIRS with capillary refill compared to venous occlusion is unclear, but a mismatch in the path length model in this volunteer would appear to be an unlikely explanation, as this would have also led to poor venous occlusion data. One feasible explanation could be that there was a pressure induced change in the tissue optical properties, for example movement of interstitial fluid or compression of the tissue followed by slow relaxation once pressure was released.

Healthy adult volunteer 4

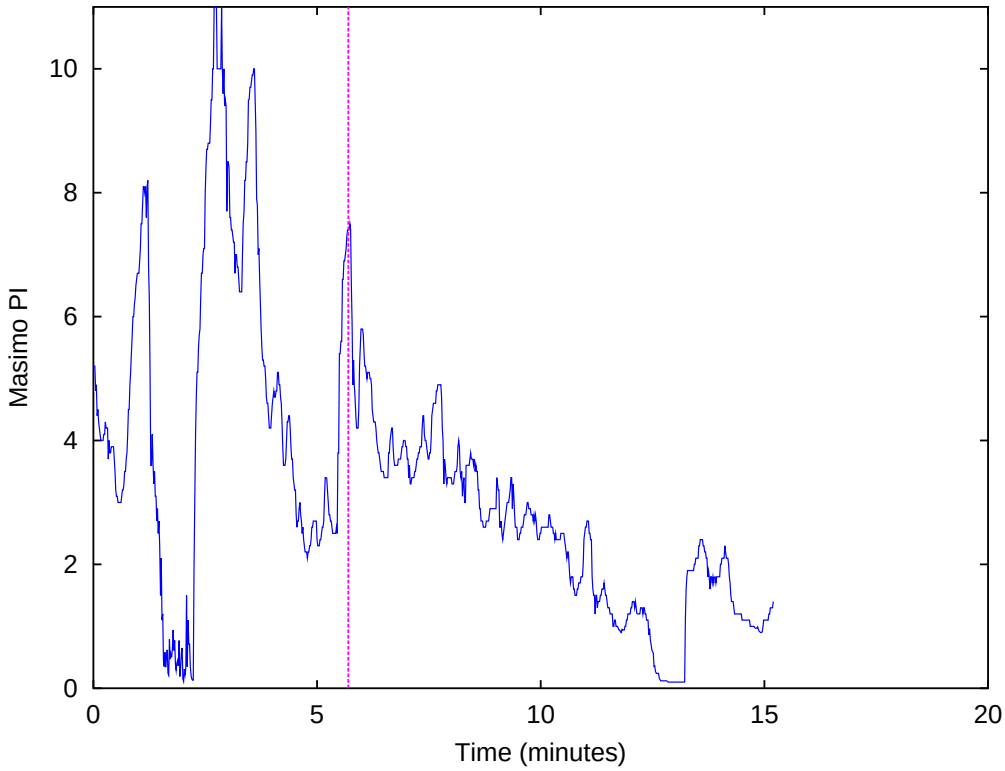


Figure 10.78: Masimo PI recorded from finger of the fourth healthy adult volunteer. There may be some evidence of reactive hyperaemia in this volunteer from the spikes in PI metric following the arterial occlusions. The water bath was at 9°C.

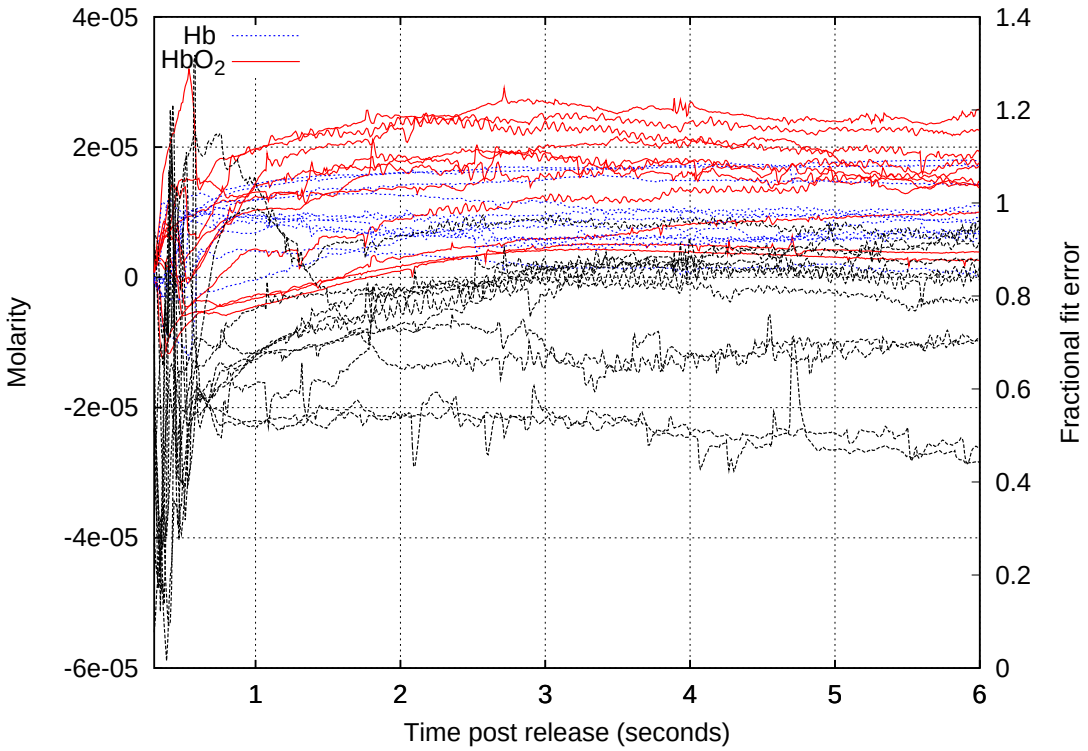


Figure 10.79: Raw NIRS based molarity recorded during room temperature capillary refills, with and without arterial occlusion.

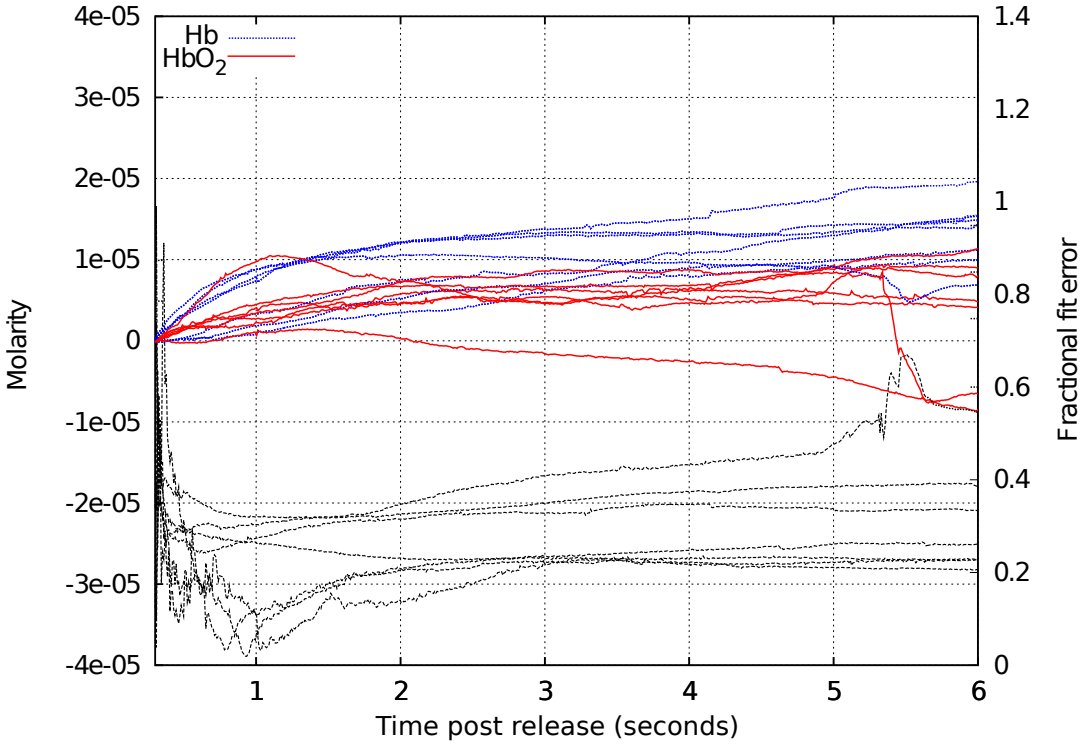


Figure 10.80: As above, recorded during post water bath capillary refills.

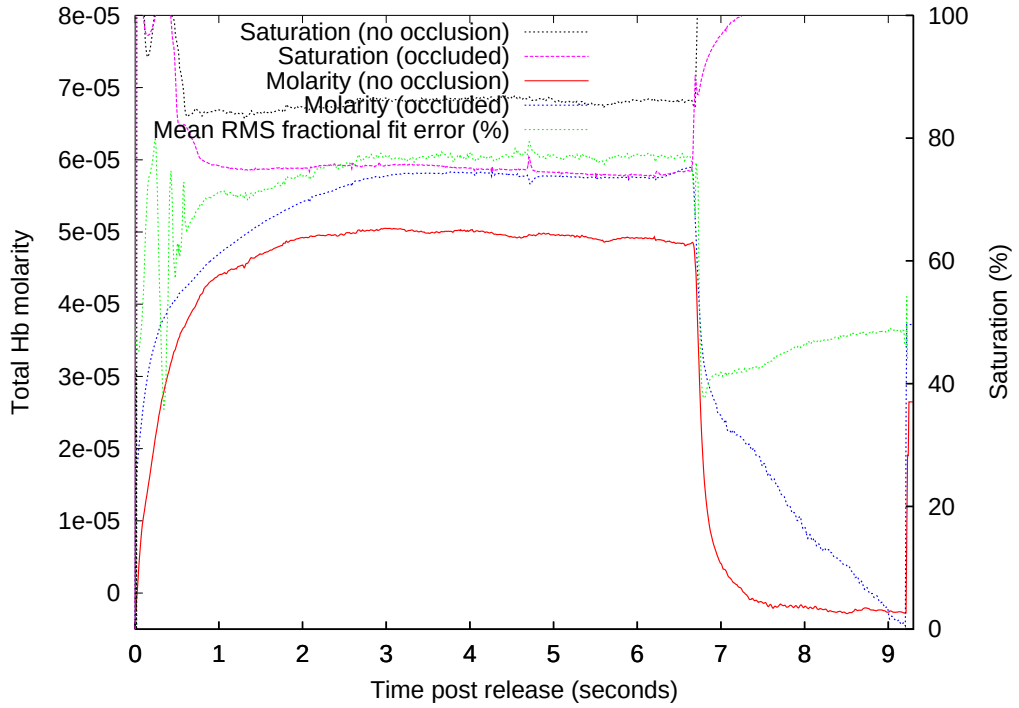


Figure 10.81: Weighted mean NIRS based molarity, error and saturation recorded during room temperature capillary refills, with and without arterial occlusion.

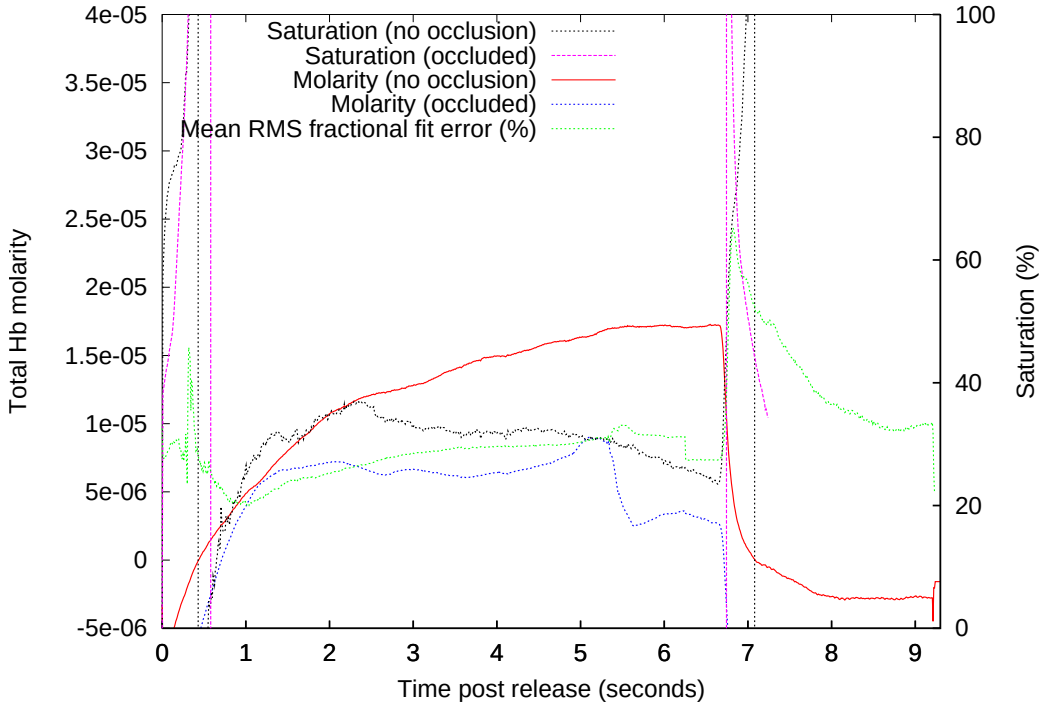


Figure 10.82: As above, recorded during post water bath capillary refills.

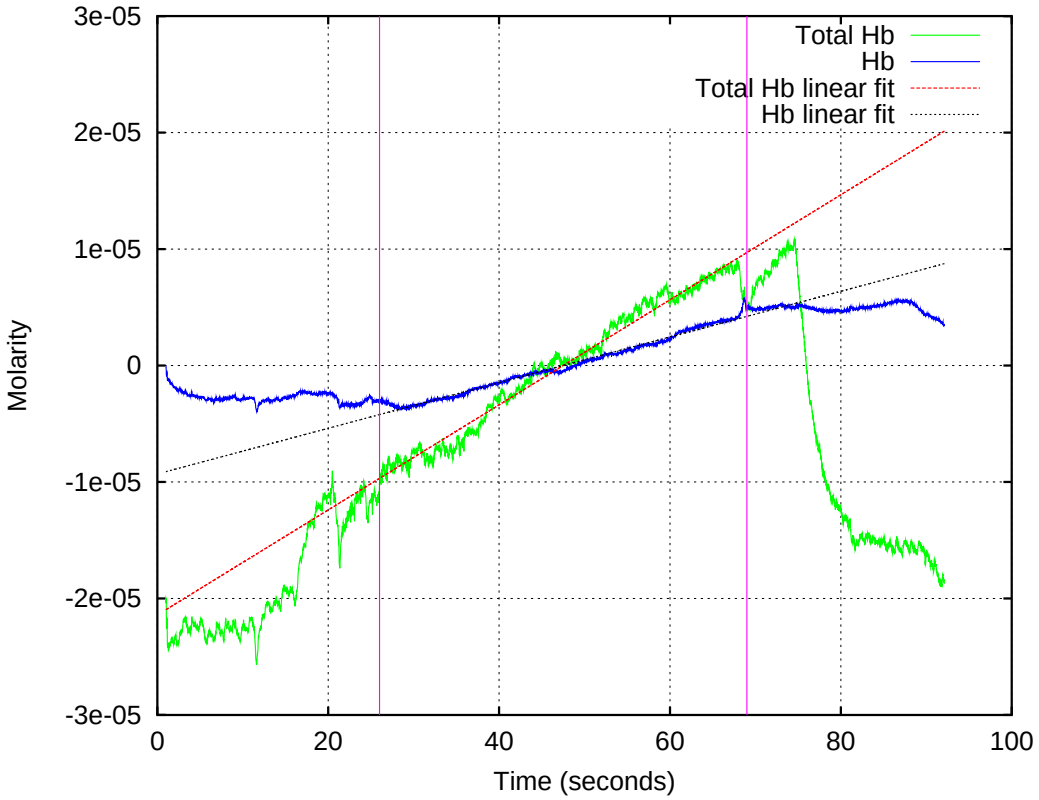


Figure 10.83: Venous occlusion NIRS recorded at room temperature.

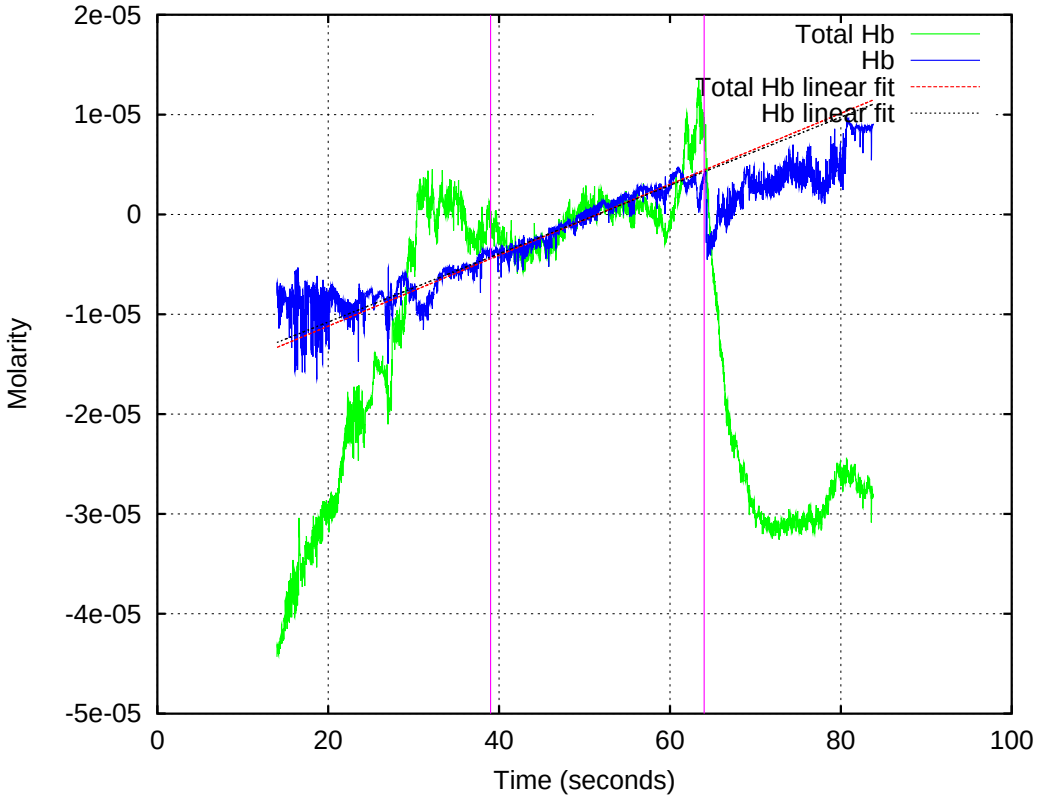


Figure 10.84: As above, post water bath.

The capillary refill data from this volunteer were of low quality, with the post water bath refills having a very high fit error, and the room temperature refills having significant coupling induced artefacts. The reasons for this are unclear, as manual observation of the sensor indicated that it had good skin contact during all the refills.

Experiment	Perfusion (ml/100g/min)	Venous saturation (%)	Metabolic rate (mlO_2/100g/min)
Room temperature	1.18 ±0.26	56.5 ±11.7	0.120 ±0.009
Water bath cooled	0.72 ±0.37	3.7 ±34.0	0.185 ±0.031

Table 10.14: Fit results and errors from the venous occlusion/NIRS.

The venous occlusion data indicate a perfusion decrease during the water bath exposure, but this decrease was slightly below the threshold for statistical significance. There was a statistically significant change in the metabolic rate, which would not be expected, so it would appear that the accuracy of the NIRS in this volunteer is in doubt. However, both metabolic rate values are within the range found in previous studies.

Conclusions

To further investigate the effect of arterial occlusion and environmental temperature changes on capillary refill in the volunteers, the "effective refill perfusion" was calculated for each of the volunteers (see section 10.6.1 for further details). In figure 10.85, the weighted mean of this value (using the weighting procedure outlined in the previous section) is plotted together with vertical error bars estimated using equation 10.28. For several of the volunteers this process involved as few as 3 or 4 capillary refills, as it was necessary to reject refills where there was either serious sensor contact variability (visible as large baseline intensity changes) or the pressure cuff was still inflating. From section 10.3, it is clear that an error estimate made with so few data-points is unlikely to be accurate, so the individual refills are also indicated using crosses. It can be seen that cooling decreased the refill perfusion in all 4 volunteers, with arterial occlusion lowering perfusion further in three volunteers. Effective perfusion during refill would appear to be sensitive to true perfusion, but the sensitivity is highly variable, and there is high inter subject baseline variance. These findings mirror the conclusions of the PaedCEW study; that capillary refill is only a useful metric of perfusion change, and has a high inter subject variability in its sensitivity and baseline value.

The data from the Masimo Radical 7 device indicate that PI provides a measurement of perfusion: PI fell with cooling in all volunteers, and was close to zero when arterial occlusion was applied. PI was vulnerable to limb motion, leading to significant artefacts

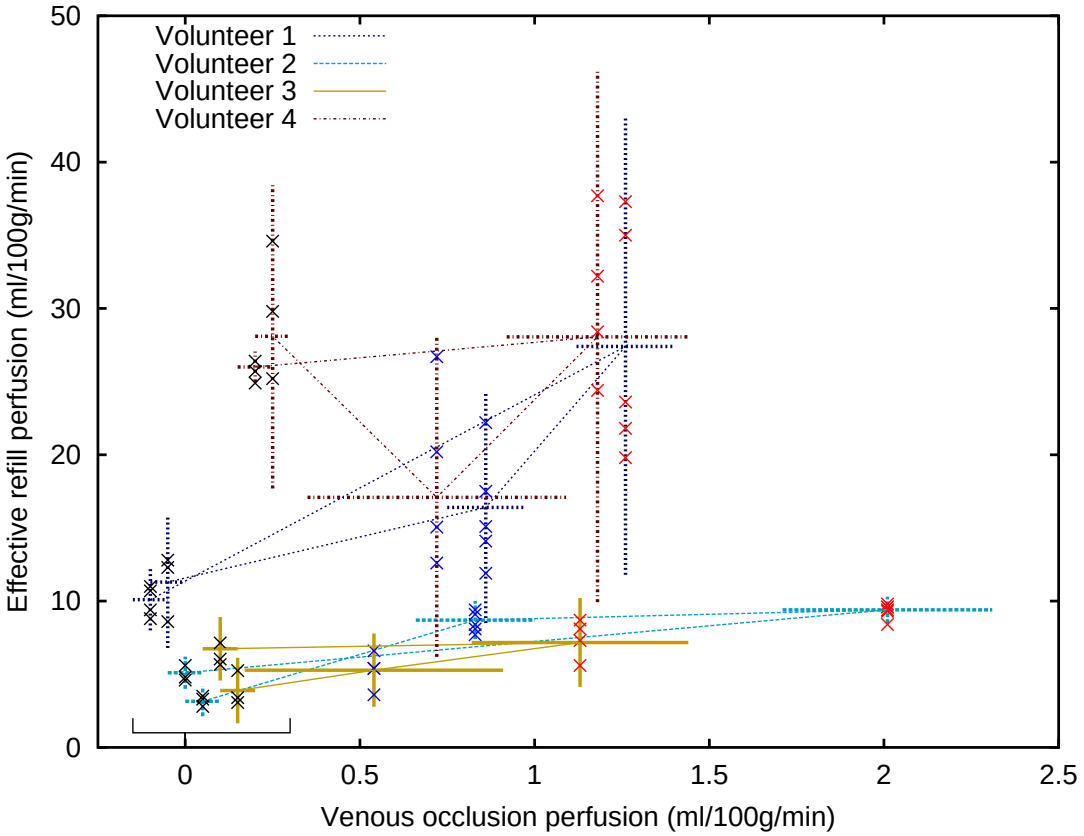


Figure 10.85: Venous occlusion perfusion compared with the effective perfusion during refill, data from the four healthy adult volunteers.

in some of the recordings. Unlike the MoorVMS laser Doppler there is no means to determine if motion is present, although the signal integrity functionality built into the device appears to filter out the effects of severe motion (section 10.4.2). Note that the pulse oximeter uses a finger clip sensor, so the sensed perfusion is not at the same site as the NIRS sensor.

To further investigate the relationship between PI and true perfusion, mean PI was calculated for the time periods prior to and post water bath immersion, and for the duration of each arterial occlusion. These values were plotted against venous occlusion perfusion for each volunteer. A very approximated vertical error is shown for each period of averaged PI, with the bars extending to the quartiles.

In figure 10.86, to the tolerances allowed by the errors bars, there is a linear relationship between PI and true volumetric perfusion in all volunteers, with zero offset at the origin. However, the gradient differs substantially between volunteers, varying from approximately 1.25 to 4 $(PIunits)/(ml/100g/min)$, with a mean value of $\sim$2. To further investigate whether this is an inter or intra subject variability, a second water bath experiment was conducted with volunteer 2, this time measuring only venous occlusion and PI (raw data not shown for succinctness). Although this repeated experiment saw

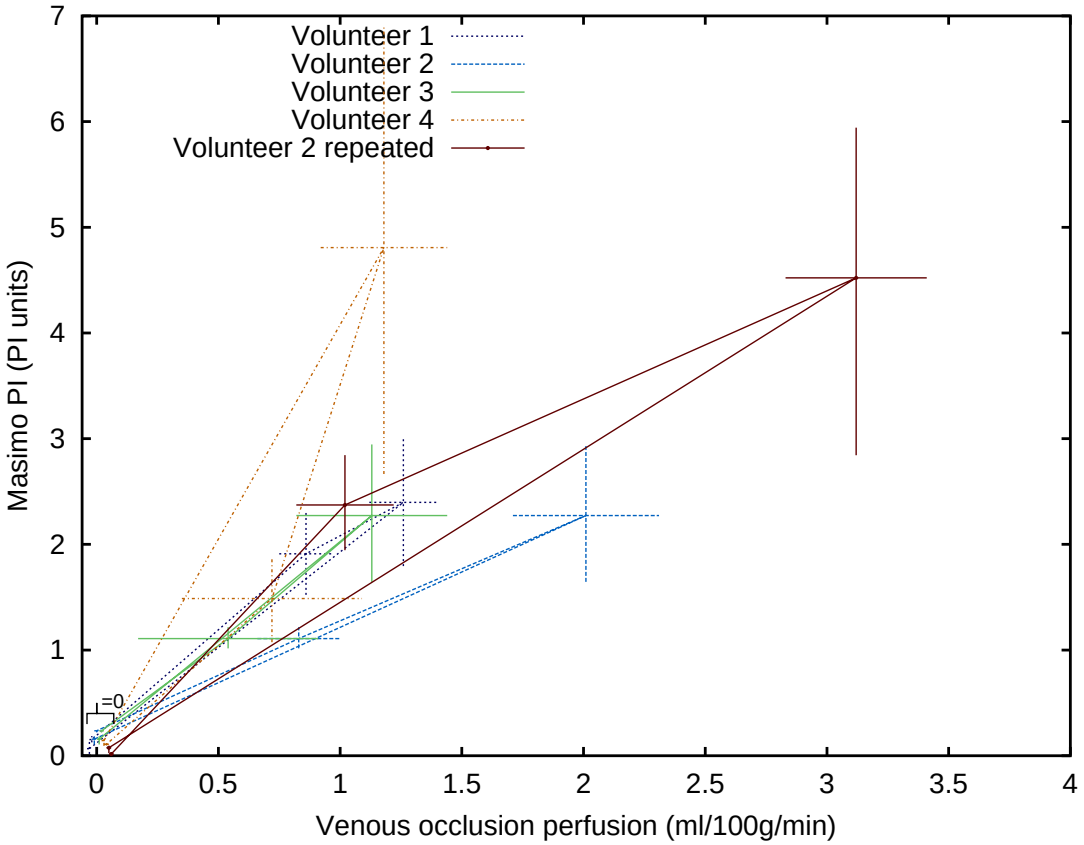


Figure 10.86: Masimo PI compared with the effective perfusion during refill, data from the four healthy adult volunteers. An offset has been added at the origin for clarity, denoted by "=0".

a higher gradient, it is unclear how significant the difference is. It would appear that the gradient differences between the volunteers are likely to be significant, but this not been proven with a formal statistical method either. However, data from the PaedCEW study show a significant shift in baseline PI level following sensor reattachment in volunteers V7C, V13C, V14C, V16C, V8A and V12A (see section 10.5.2), strong evidence that a large component of the sensitivity variability arises from sensor attachment.

Although PI fell to close to zero under arterial occlusion, it is unclear if occlusion realistically simulates very low perfusion, for example in the case of severe vasoconstriction leading to close to zero perfusion, there may still be arterial pulsation potentially giving rise to a non zero PI value. Discarding the occlusion data, it is not possible to produce accurate estimates of the y axis intercept from analysis of the prior and post water bath immersion periods only, but a repeated experiment using a larger group of volunteers together with a more accurate NIRS device (for example an off the shelf time domain device) may allow accurate estimation of both the distribution of gradients (i.e. PI sensitivity to true perfusion), and the y intercept point (i.e. the baseline PI in an unperfused state where arterial blood is still pulsatile) without the use of arterial occlusion.

Chapter 11

Conclusions and potential future research

11.1 Summary and conclusions

The majority of the experimental work in this thesis has been directed towards the development of Capillary Refill Timing (CRT). Although CRT is now a more mature technology, with a fully automated portable device developed in Chapter 9, results from the PaedCEW study indicate that capillary refill time is not as effective as reflection mode laser Doppler at detecting a cooling induced perfusion change. However, the extent to which the forearm cooling resulted in a decrease in perfusion is unclear, as there was no "gold standard" perfusion metric in use as part of the study.

In the PaedCEW study, a statistically significant correlation was seen between short period (<5 minute period, and comprised principally of ∼10 second period Mayer waves) oscillations in the laser Doppler metric (Perfusion Units) and capillary refill time. Experimental results (Chapter 10) from the NIRS hardware developed in Chapter 8 support the conclusion that following blanching, capillaries principally refill with blood from adjacent tissue through the micro-vasculature. The PaedCEW study does not demonstrate any direct measurement of volumetric perfusion via capillary refill; if the theoretical explanation developed in Chapter 10 is an accurate one, it would appear that the capillary refill time metric is instead a measure of some function of Peripheral Vascular Resistance (PVR) and perfusion pressure (MAP-MVP, i.e. the mean pressure drop across the capillary bed). As volumetric perfusion is also a function of perfusion pressure and PVR, a correlation between refill time and volumetric perfusion is to be expected.

In conclusion, it would appear from the study data that, although the prototype capillary refill sensor was found to be less reliable than the MoorVMS laser Doppler device at detecting a temperature induced perfusion change, capillary refill time provides a

useful measure of perfusion trend, similar in utility to laser Doppler PU, with better performance than pulse oximetry based Perfusion Index (PI) in the volunteer study, and potentially with greater resistance to motion induced signal artefacts than laser Doppler. Capillary refill time significantly outperformed PI in the child volunteers, but PI performance increased in adult volunteers, and was highly correlated with volumetric perfusion calculated using venous occlusion/NIRS in a small scale study of adult volunteers. This may be indicative of a physiological difference between children and adults, but a repeat of this experiment with child volunteers is needed to make conclusions. The low production cost of a capillary refill sensor may be significant in considering its utility; the MoorVMS probes used in the PaedCEW study were multiple use devices costing approximately $1k each, likely due to their single mode fibres and an electrical/optical connector to the base unit, whereas the pneumatic refill sensor prototypes discussed in Chapter 9 could be mass produced as disposable devices using low cost materials.

11.2 Suggested future work

In this thesis, a diverse range of non-invasive peripheral perfusion monitoring techniques have been explored to varying degrees, summarised in table 11.1. In the table, the "type of metric" column header refers to the relationship between perfusion and the measured quantity or property, and Technology Readiness Level (TRL) to the current development state. A detailed explanation of TRL can be found in a white paper by Mankins which outlines the concept [270], and the precise reasoning behind the value assigned to each technique can be found in Chapter 7.

The two final columns refer to the Chapters in which each technique was explored further, the first referring to quantitative studies of theoretical operation, i.e. both numerical simulation and experimental study not on human volunteers, and the second to representative tests on healthy volunteers, i.e. both qualitative and quantitative study in a relevant environment. Healthy volunteer studies comprised both the PaedCEW study (conducted jointly with the medical school, using CRT, Masimo perfusion index and reflection mode laser Doppler with a MoorVMS sensor), and a set of experiments designed to investigate the origin of the refilling blood in the capillary refill test, which made use of simultaneous Masimo PI, venous occlusion with NIRS, and capillary refill with NIRS.

The PaedCEW study involved the use of an automated capillary refill device on healthy adult and child volunteers. Automated capillary refill timing has not been demonstrated in previously published studies, meaning the PaedCEW study raised the technology readiness level from 3; analytical and experimental critical function and/or characteristic

Method	Type of metric	TRL	Theoretical studies presented	Data presented from healthy volunteers
Thermal: active	Absolute	4-6	Chapter 2	No
Thermal: passive	Absolute	3	No	No
NMR (permanent magnet)	Absolute	3	Chapter 3	No
Trans-cutaneous blood gas: O_2 and CO_2	Adequacy	7-9	No	No
Trans-cutaneous blood gas: H_2 washout	Absolute	5	No	No
Ultrasonic: Colour Doppler	Absolute theoretically achievable	7	No	No
Ultrasonic: Photo-acoustic	Unclear, potentially absolute	6	No	No
Laser Doppler: Reflection mode QELS	Relative: Hight inter and intra subject baseline variability	9	No	Chapter 10
Laser Doppler: DWS/DCS	Relative: Potential reduction in inter and intra subject baseline variability	7	No	No
NIRS: tissue oxygenation	Adequacy	9	No	No
NIRS: venous occlusion	Absolute	7	No	Chapter 10
Capillaroscopy	Relative: trend information for a given site	7-9	No	No
PPG PI (Masimo Inc.)	Relative: trend information for a given site Some degree of absolute is claimed	9	No	Chapter 10
Capillary refill time	Relative: trend information for a given site	3 (now 4-5)	Chapter 10	Chapter 10

Table 11.1: Summary of the various non-invasive perfusion monitoring methods discussed in this thesis.

proof of concept (automated capillary refill timing has been previously demonstrated with semi automated devices in small scale laboratory experiments) to 5; component and/or breadboard validation in relevant environment, as the fully automated and self contained capillary refill hardware described in Chapter 9 was designed to comply with all relevant medical device standards for use in the clinical environment, and was tested on a significant number of child and adult volunteers in the presence of changing ambient temperature and limb motion. However, neither the pneumatic data logger nor the sensor are representative of a final system for clinical use.

The TRL column shows a range of levels for some techniques (e.g. 4 to 6 for active thermal). This is due to the range of potential implementations that could be utilised; for example, active thermal sensors could using a simple heated disk, with a TRL of 6 due to published demonstration of prototype hardware in a relevant environment (by Toumi et al[25]), or alternatively, a novel implementation using an inner disc and outer ring, as discussed in Chapter 2, demonstrated only in numerical simulation work could be assigned TRL 4 (although this simulation was based on well understood principles).

Similarly for trans-cutaneous blood gas sensing, a TRL range of 7 to 9 was assigned. Trans-cutaneous O_2 and CO_2 sensors are in widespread clinical use, and there is evidence to support their use in the assessment of perfusion adequacy, so such an implementation using existing sensors could be categorised as TRL 9. However, there is also evidence (Chapter 4) that unconventional operating modes such as use of CO_2 sensing without skin heating may give improved performance as a perfusion adequacy metric. Such use has not been demonstrated in regular clinical practice, so a TRL of 7 could been assigned, meaning there is potential for useful future studies to be made into the development of unheated trans-cutaneous CO_2 sensing as a perfusion adequacy metric, although this would be unlikely to involve development of novel hardware.

Capillaroscopy was assigned a TRL of 7-9, based on the fact that nailbed capillaroscopy is an off-the-shelf technology, with commercial devices able to provide perfusion metrics such as functional capillary density. Such an implementation could be assigned a TRL of 9, although it should be noted that this would be an impractical solution for continuous monitoring, and does not provide a very direct measurement of perfusion. Custom made compact sensors, and erythrocyte velocimetry via machine vision have both been demonstrated, and may allow for a more effective capillascope based perfusion sensor, but these techniques have a considerably lower readiness level.

It should be noted that the readiness level does not necessarily relate to the likelihood that a given technology could be developed into a successful medical device, only to how far along to pathway to development it is at the present time. For example, both

NMR and passive thermal monitoring have a TRL of 3, meaning that analytical proof of concept exists for the core principles of operation. However, a detailed analysis of the practicality of NMR, conducted in Chapter 3, led to the conclusion that a permanent magnet based NMR perfusion sensor is unlikely to be a practical device. On the other hand, simple forms of passive thermal perfusion monitoring such as core to periphery temperature difference are already in use, meaning there is precedent for the core principles. The basic theoretical understanding of the processes involved in heat transport from the perfusing blood is also well developed, so for this reason it would appear that a passive thermal perfusion sensor based on the measurement of skin surface thermal flux (as discussed in Chapter 2) has a high probability of leading to a successful device.

Another important consideration is the current level of commercial interest in the different techniques discussed in this thesis; the majority of the methods with a TRL in the range of 7 to 9 are already in relatively widespread use, with a large range of off-the-shelf devices available, and considerable invested development effort. For these reasons significant improvements may be unlikely in trans-cutaneous O_2 or CO_2 sensor design, NIRS, Doppler ultrasound (colour Doppler), or PPG based perfusion index. However, there may be significant opportunities for future improvements in the application of capillaroscopy and venous occlusion with NIRS as perfusion metrics, as these techniques have not been developed with routine perfusion sensing as a design goal.

Following an initial literature review and exploration of potential monitoring methods, in Chapter 7 optical perfusion monitoring methods were selected as the most promising route for experimental study. However, several of the methods that were not explored further have clear potential for future development. These are:

- Thermal perfusion monitoring: recent published studies have shown good correlation between laser Doppler and active thermal [25], and numerical simulations (Chapter 2) demonstrate the potential of novel sensor geometries. The simulation studies could be expanded upon by construction of sensor prototypes followed by studies on tissue phantoms and healthy volunteers.

- Trans-cutaneous H_2 clearance: other than studies by Harrison et al [7, 271], this method has seen little development. Fully trans-cutaneous H_2 clearance sensing, with H_2 being introduced via diffusion through the skin, would be a necessary development for this method to become practical, but has not been demonstrated in published work. Development of such a prototype trans-cutaneous sensor would allow the performance of this method to be studied.

- Perfusion adequacy assessment via “Discrete Saturation Transform” (US Patent

US5632272 A): this was briefly discussed in Chapter 6, but has not been studied in detail as part of this thesis. Potentially this signal processing technique may allow the venous saturation to be derived using a conventional pulse oximetry sensor, although it should be noted that venous saturation is not necessarily a measure of perfusion adequacy, as arterial to venous shunts allow blood to flow directly from arteries to veins, bypassing the capillary bed. Studies on healthy volunteers using a commercial pulse oximetry sensor interfaced with the data logger system developed in Chapter 9 would allow for an assessment of this method to be made.

Although the experimental studies discussed in Chapter 10 have allowed conclusions to be drawn regarding the nature of the capillary refill, there is clear room for improvement if these were to be repeated. The pressure application device (development of which was discussed in Chapter 8) could be improved by replacement of the stepper motor with a servo motor driven by an off-the-shelf control unit operating in constant torque mode, allowing more accurate and repeatable pressure applications to be made. Measurement accuracy could also be improved through use of an off-the-shelf NIRS device in the venous occlusion experiments.

Repeating the PaedCEW study with a means of direct perfusion measurement would allow the laser Doppler, perfusion index and capillary refill time metrics to be directly related to true perfusion, rather than merely to each other and skin surface temperature. The influence of sensor attachment on CRT variability could also be studied.

The pneumatic capillary refill device developed in Chapter 9 has set the groundwork for future CRT sensor developments. The PaedCEW study used a somewhat large and cumbersome sensor described in Appendix A, and it appears likely that there was considerable capillary refill time biassing due to the large thermal mass and insulating properties of this sensor. Use of off-the-shelf blood pressure cuffs as pressure applicators placed a lower limit on sensor size, and considerably increased its cost.

Development of a prototype vacuum formed sensor suitable for mass production is discussed in Chapter 9. Unfortunately, due to difficulties in reliably forming the polyurethane outer cover and adhering the components parts of the outer cover to form an airtight seal, it did not prove possible to develop this design to the point where it was suitable for use in the PaedCEW study. Attention from an expert in polymer manufacturing is required before further progress can be made.

Along with an improved capillary refill sensor, the author would also recommend the development of an active thermal perfusion sensor, based on the design of Toumi et al, but incorporating the dual element construction simulated in Chapter 2.

Bibliography

[1] Lima, Alexandre and Jansen, Tim C and van Bommel, Jasper and Ince, Can and Bakker, Jan, "The prognostic value of the subjective assessment of peripheral perfusion in critically ill patients," *Critical care medicine*, vol. 37, no. 3, pp. 934–938, 2009.

[2] Boyd, O and Bennett, ED, "Enhancement of perioperative tissue perfusion as a therapeutic strategy for major surgery.," *New horizons (Baltimore, Md.)*, vol. 4, no. 4, pp. 453–465, 1996.

[3] Shoemaker, William C and Thangathurai, Duraiyah and Wo, Charles CJ and Kuchta, Kenneth and Canas, Marcos and Sullivan, Michael J and Farlo, Joseph and Roffey, Peter and Zellman, Vladimir and Katz, Ronald L, "Intraoperative evaluation of tissue perfusion in high-risk patients by invasive and noninvasive hemodynamic monitoring," *Critical care medicine*, vol. 27, no. 10, pp. 2147–2152, 1999.

[4] Furnas, Heather and Rosen, Joseph M, "Monitoring in microvascular surgery," *Annals of plastic surgery*, vol. 26, no. 3, pp. 265–272, 1991.

[5] Abdel-Galil, Khalid and Mitchell, David, "Postoperative monitoring of microsurgical free-tissue transfers for head and neck reconstruction: a systematic review of current techniques-part II. Invasive techniques," *British Journal of Oral and Maxillofacial Surgery*, vol. 47, no. 6, pp. 438–442, 2009.

[6] Hardwicke, Joseph and Thomson, Richard and Bamford, Amy and Moiemen, Naiem, "A pilot evaluation study of high resolution digital thermal imaging in the assessment of burn depth," *Burns*, vol. 39, no. 1, pp. 76–81, 2013.

[7] Harrison, DK and Raad, R Abi and Newton, D and McCollum, PT, "Transcutaneous hydrogen clearance-a new non-invasive technique for assessing blood flow in human skin," *Physiological measurement*, vol. 15, no. 1, p. 89, 1994.

[8] Tatevossian, Raymond G and Wo, Charles CJ and Velmahos, George C and Demetriades, Demetrios and Shoemaker, William C, "Transcutaneous oxygen and CO2 as early warning of tissue hypoxia and hemodynamic shock in critically ill emergency patients," *Critical care medicine*, vol. 28, no. 7, pp. 2248–2253, 2000.

[9] Boldt, Joachim, "Clinical review: Hemodynamic monitoring in the intensive care unit," *Critical Care*, vol. 6, no. 1, p. 52, 2002.

[10] Baker, Roger C, *Flow measurement handbook: industrial designs, operating principles, performance, and applications*. Cambridge University Press, 2005.

[11] Kutz, Myer and Adrezin, Ronald S and Barr, Ronald E and Batich, Christopher and Bellamkonda, Ravi V and Brammer, Anthony J and Buchanan, Thomas S and Cook, Albert M and Currie, John Michael and Dolan, Alfred M and others, *Standard handbook of biomedical engineering and design*. McGraw-Hill, 2003.

[12] Tibby, Shane M and Hatherill, Mark and Murdoch, Ian A, "Capillary refill and core–peripheral temperature gap as indicators of haemodynamic status in paediatric intensive care patients," *Archives of disease in childhood*, vol. 80, no. 2, pp. 163–166, 1999.

[13] Osborn, DA and Evans, N and Kluckow, M, "Clinical detection of low upper body blood flow in very premature infants using blood pressure, capillary refill time, and central-peripheral temperature difference," *Archives of Disease in Childhood-Fetal and Neonatal Edition*, vol. 89, no. 2, pp. F168–F173, 2004.

[14] Lima, Alexandre Pinto and Beelen, Peter and Bakker, Jan, "Use of a peripheral perfusion index derived from the pulse oximetry signal as a noninvasive indicator of perfusion," *Critical care medicine*, vol. 30, no. 6, pp. 1210–1213, 2002.

[15] Kelly, Greg, "Body temperature variability (Part 1): a review of the history of body temperature and its variability due to site selection, biological rhythms, fitness, and aging," *Alternative medicine review*, vol. 11, no. 4, pp. 278–293, 2006.

[16] "Hemedex Inc. website." `http://www.hemedex.com/Products_Page.htm`, June 2011.

[17] Lv, Yong-Gang and Liu, Jing, "Measurement of local tissue perfusion through a minimally invasive heating bead," *Heat and Mass Transfer*, vol. 44, no. 2, pp. 201–211, 2007.

[18] Robinson, Paul S and Scott, Elaine P and Diller, Thomas E, "Testing of a Non-invasive Probe for Measurement of Blood Perfusion," *Journal of medical devices*, vol. 2, no. 1, 2008.

[19] Li, Vivian Victoria, "Design of a thermal diffusion sensor for noninvasive assessment of skin surface perfusion and endothelial dysfunction," Master's thesis, Massachusetts Institute of Technology, 2008.

[20] Liu, Jing and Xu, Lisa X, "Estimation of blood perfusion using phase shift in temperature response to sinusoidal heating at the skin surface," *Biomedical Engineering, IEEE Transactions on*, vol. 46, no. 9, pp. 1037–1043, 1999.

[21] Song, Chang W, "Effect of local hyperthermia on blood flow and microenvironment: a review," *Cancer Research*, vol. 44, no. 10 Supplement, pp. 4721s–4730s, 1984.

[22] Cardinali, Alex V, "Validation of a Noninvasive Blood Perfusion Measurement Sensor," Master's thesis, Virginia Polytechnic Institute and State University, 2002.

[23] Toumi, D. and Gehin, C. and Dittmar, A. and McAdams, E., "Design and validation of an ambulatory system for the measurement of the microcirculation in the capillaries: the Hematron device," in *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE*, pp. 4120–4123, Sept 2009.

[24] Dittmar, A and Pauchard, T and Delhomme, G and Vernet-Maury, E, "A thermal conductivity sensor for the measurement of skin blood flow," *Sensors and actuators B: Chemical*, vol. 7, no. 1, pp. 327–331, 1992.

[25] Toumi, D and Gehin, C and Bigouret, A and Dittmar, A and McAdams, ET, "Comparison between ambulatory measurement of effective thermal conductivity and laser Doppler flowmetry method to assess skin microcirculatory activity," *IRBM*, 2014.

[26] Bowman, Harry F and Weaver, James C, "Apparatus for the non-invasive measurement of thermal properties and perfusion rates of biomaterials," Aug 22 1989. US Patent 4,859,078.

[27] Pennes, HH, "Analysis of skin, muscle and brachial arterial blood temperatures in the resting normal human forearm.," *The American journal of the medical sciences*, vol. 215, no. 3, p. 354, 1948.

[28] Rawnsley, RJ and Roemer, RB and Dutton, AW and others, "The simulation of discrete vessel effects in experimental hyperthermia.," *Journal of biomechanical engineering*, vol. 116, no. 3, p. 256, 1994.

[29] Arkin, H and Xu, LX and Holmes, KR, "Recent developments in modeling heat transfer in blood perfused tissues," *Biomedical Engineering, IEEE Transactions on*, vol. 41, no. 2, pp. 97–107, 1994.

[30] Brinck, Heinrich and Werner, Jijrgen, "Efficiency function: improvement of classical bioheat approach," *Journal of Applied Physiology*, vol. 77, no. 4, pp. 1617–1622, 1994.

[31] Wissler, Eugene H, "Pennes' 1948 paper revisited," *Journal of Applied Physiology*, vol. 85, no. 1, pp. 35–41, 1998.

[32] C. CK., "Mathematical models of bioheat transfer.," vol. 22:19, p. 155, 1992.

[33] Jiji, Latif Menashi, *Heat conduction*. Springer Berlin, 2009.

[34] Wulff, W, "The energy conservation equation for living tissue," *Biomedical Engineering, IEEE Transactions on*, no. 6, pp. 494–495, 1974.

[35] Klinger, HG, "Heat transfer in perfused biological tissue-I: General theory," *Bulletin of mathematical biology*, vol. 36, pp. 403–415, 1974.

[36] Chen, Michael M and Holmes, Kenneth R, "Microvascular contributions in tissue heat transfer," *Annals of the New York Academy of Sciences*, vol. 335, no. 1, pp. 137–150, 1980.

[37] Kreith, F and Timmerhaus, K and Lior, N and Shaw, H and Shah, RK and Bell, KJ and others, "The CRC Handbook of Thermal Engineering. Ed. Frank Kreith Boca Raton: CRC Press LLC, 2000," 2000.

[38] Weinbaum, SJLM and Jiji, LM and Lemons, DE, "Theory and experiment for the effect of vascular microstructure on surface tissue heat transfer–Part I: Anatomical foundation and model conceptualization.," *Journal of biomechanical engineering*, vol. 106, no. 4, p. 321, 1984.

[39] Jiji, LM and Weinbaum, S and Lemons, DE, "Theory and experiment for the effect of vascular microstructure on surface tissue heat transfer–Part II: Model formulation and solution.," *Journal of biomechanical engineering*, vol. 106, no. 4, p. 331, 1984.

[40] Weinbaum, S and Jiji, LM, "A two phase theory for the influence of circulation on the heat transfer in surface tissue," *Advances in Boiengineering, ASME, New York*, pp. 179–182, 1979.

[41] Weinbaum, S and Jiji, LM and others, "A new simplified bioheat equation for the effect of blood flow on local average tissue temperature.," *Journal of biomechanical engineering*, vol. 107, no. 2, p. 131, 1985.

[42] Wren, M. Karlsson and D. Loyd, J., "A hybrid equation for simulation of perfused tissue during thermal treatment," *International Journal of Hyperthermia*, vol. 17, no. 6, pp. 483–498, 2001.

[43] Bowman, H Frederick and Cravalho, Ernest G and Woods, Monty, "Theory, measurement, and application of thermal properties of biomaterials," *Annual review of biophysics and bioengineering*, vol. 4, no. 1, pp. 43–80, 1975.

[44] Potocki, JK and Tharp, HS, "Concurrent hyperthermia estimation schemes based on extended Kalman filtering and reduced-order modelling," *International journal of hyperthermia*, vol. 9, no. 6, pp. 849–865, 1993.

[45] "Texas Instruments website." `http://www.ti.com/product/tmp006`, December 2012.

[46] "Melexis Semiconductor website." `http://www.melexis.com/product/mlx90616`, January 2013.

[47] Sarelius IH., "Cell flow path influences transit time through striated muscle capillaries.," *American J Physiol Heart Circ Physiol*, vol. 250, pp. H899–H907, 1986.

[48] Andrew E. Marble and Igor V. Mastikhin and Bruce G. Colpitts and Bruce J. Balcom, "A compact permanent magnet array with a remote homogeneous field," *Journal of Magnetic Resonance*, vol. 186, no. 1, pp. 100 – 104, 2007.

[49] Fagrell, BENGT and Fronek, ARNOST and Intaglietta, MARCOS, "A microscope-television system for studying flow velocity in human skin capillaries," *Am. J. Physiol*, vol. 233, no. 2, pp. H318–H321, 1977.

[50] Szafer, Aaron and Zhong, Jianhui and Gore, John C, "Theoretical model for water diffusion in tissues," *Magnetic Resonance in Medicine*, vol. 33, no. 5, pp. 697–712, 1995.

[51] K & J magnetics Inc., "N52 magnets." `http://www.kjmagnetics.com/proddetail.asp?prod=D64-N52`, November 2010.

[52] Elliott, R. S., *Electromagnetics*. IEEE Press, 1993.

[53] B. Manz, A. Coy, R. Dykstra, C. Eccles, M. Hunter, B. Parkinson, and P. Callaghan, "A mobile one-sided NMR sensor with a homogeneous magnetic field: The NMR-MOLE," *Journal of Magnetic Resonance*, vol. 183, no. 1, pp. 25 – 31, 2006.

[54] Perlo, Juan and Casanova, Federico and Blümich, Bernhard, "Ex Situ NMR in Highly Homogeneous Fields: 1H Spectroscopy," *Science*, vol. 315, no. 5815, pp. 1110–1112, 2007.

[55] Blank, Aharon and Alexandrowicz, Gil and Muchnik, Lev and Tidhar, Gil and Schneiderman, Jacob and Virmani, Renu and Golan, Erez, "Miniature self-contained intravascular magnetic resonance (IVMI) probe for clinical applications.," *Magn Reson Med*, vol. 54, no. 1, pp. 105–12, 2005.

[56] Raich, H and Blümler, P, "Design and construction of a dipolar Halbach array with a homogeneous field from identical bar magnets: NMR Mandhalas," *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering*, vol. 23, no. 1, pp. 16–25, 2004.

[57] Battocletti, Joseph H and Halbach, Richard E and Antonich, Frederick J and Sances Jr, Anthony and Knox, Thomas A, "Nuclear magnetic resonance blood flowmeter," Sept. 23 1986. US Patent 4,613,818.

[58] Battocletti, Joseph H and Linehan, John H and Larson, Sanford J and Sances, Anthony and Bowman, Robert L and Kudravcev, Vsevelod and Genthe, William K and Halbach, Richard E and Evans, Silas M, "Analysis of a nuclear magnetic resonance blood flowmeter for pulsatile flow," *Biomedical Engineering, IEEE Transactions on*, no. 6, pp. 403–407, 1972.

[59] Battocletti, Joseph H and Halbach, Richard E and Salles-Cunha, Sergio X and Sances Jr, Anthony, "The NMR blood flowmeter—theory and history," *Medical Physics*, vol. 8, p. 435, 1981.

[60] Appelt, Stephan and Kühn, Holger and Häsing, F Wolfgang and Blümich, Bernhard, "Chemical analysis by ultrahigh-resolution nuclear magnetic resonance in the Earth's magnetic field," *Nature Physics*, vol. 2, no. 2, pp. 105–109, 2006.

[61] Mag-Inc., "Kool-mμ." `http://www.mag-inc.com/products/powder-cores/kool-mu`, November 2010.

[62] Varghese, Oomman K and Yang, Xiping and Kendig, James and Paulose, Maggie and Zeng, Kefeng and Palmer, Charles and Ong, Keat Ghee and Grimes, Craig A,

"A transcutaneous hydrogen sensor: From design to application," *Sensor Letters*, vol. 4, no. 2, pp. 120–128, 2006.

[63] Rüdiger, Mario and Töpfer, Kerstin and Hammer, Hannes and Schmalisch, Gerd and Wauer, Roland R, "A survey of transcutaneous blood gas monitoring among European neonatal intensive care units," *BMC pediatrics*, vol. 5, no. 1, p. 30, 2005.

[64] Dear, PR, "Monitoring oxygen in the newborn: saturation or partial pressure?," *Archives of disease in childhood*, vol. 62, no. 9, pp. 879–881, 1987.

[65] Brünstler, Ingrid and Enders, Angelika and Versmold, Hans T, "Skin surface PCO_2 monitoring in newborn infants in shock: Effect of hypotension and electrode temperature," *The Journal of pediatrics*, vol. 100, no. 3, pp. 454–457, 1982.

[66] Lübbers, Dietrich W, "Theoretical basis of the transcutaneous blood gas measurements," *Critical care medicine*, vol. 9, no. 10, pp. 721–733, 1981.

[67] Cabal, Luis and Hodgman, Joan and Siassi, Bijan and Plajstek, Carolyn, "Factors affecting heated transcutaneous Po2 and unheated transcutaneous Po2 in preterm infants," *Critical Care Medicine*, vol. 9, no. 4, pp. 298–304, 1981.

[68] Yamazoe, N and Hisamoto, J and Miura, N and Kuwata, S, "Potentiometric solid-state oxygen sensor using lanthanum fluoride operative at room temperature," *Sensors and actuators*, vol. 12, no. 4, pp. 415–423, 1987.

[69] Clark, Leland C and Wolf, Richard and Granger, Donald and Taylor, Zena and others, "Continuous recording of blood oxygen tensions by polarography," *Journal of applied physiology*, vol. 6, no. 3, pp. 189–193, 1953.

[70] Lucisano, Joseph Y and Armour, Jon C and Gough, David A, "In vitro stability of an oxygen sensor," *Analytical chemistry*, vol. 59, no. 5, pp. 736–739, 1987.

[71] Mackereth, FJH, "An improved galvanic cell for determination of oxygen concentrations in fluids," *Journal of Scientific Instruments*, vol. 41, no. 1, p. 38, 1964.

[72] Eberhard, Patrick, "The design, use, and results of transcutaneous carbon dioxide analysis: current and future directions," *Anesthesia & Analgesia*, vol. 105, no. 6S Suppl, pp. S48–S52, 2007.

[73] Fog, Agner and Buck, Richard P, "Electronic semiconducting oxides as pH sensors," *Sensors and Actuators*, vol. 5, no. 2, pp. 137–146, 1984.

[74] Kinlen, Patrick J and Wagenknecht, John H, "Electronic pH and ORP indicator," June 8 1993. US Patent 5,218,304.

[75] Aukland, Knur and Bower, Bruce F and Berliner, Robert W, "Measurement of local blood flow with hydrogen gas," *Circulation Research*, vol. 14, no. 2, pp. 164–187, 1964.

[76] Pasztor, Emil and Symon, Lindsay and Dorsch, NWC and Branston, Neil M, "The hydrogen clearance method in assessment of blood flow in cortex, white matter and deep nuclei of baboons," *Stroke*, vol. 4, no. 4, pp. 556–567, 1973.

[77] Singh, S and Singh, J, "Transdermal drug delivery by passive diffusion and iontophoresis: a review," *Medicinal research reviews*, vol. 13, no. 5, pp. 569–621, 1993.

[78] Malone, JM and Leal, JM and Moore, WS and Henry, RE and Daly, MJ and Patton, DD and Childers, SJ, "The "gold standard" for amputation level selection: Xenon-133 clearance," *Journal of Surgical Research*, vol. 30, no. 5, pp. 449–455, 1981.

[79] Moser, Katrin and Kriwet, Katrin and Naik, Aarti and Kalia, Yogeshvar N and Guy, Richard H, "Passive skin penetration enhancement and its quantification in vitro," *European Journal of Pharmaceutics and Biopharmaceutics*, vol. 52, no. 2, pp. 103–112, 2001.

[80] Engel, Rolf R, "Topical agent for transcutaneous measurement of partial pressure of oxygen," Dec. 18 1984. US Patent 4,488,557.

[81] Stosseck, Klaus and Lübbers, Dietrich W and Cottin, Norbert, "Determination of local blood flow (microflow) by electrochemically generated hydrogen," *Pflügers Archiv*, vol. 348, no. 3, pp. 225–238, 1974.

[82] Robergs, Robert A and Ghiasvand, Farzenah and Parker, Daryl, "Biochemistry of exercise-induced metabolic acidosis," *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, vol. 287, no. 3, pp. R502–R516, 2004.

[83] Lima, Alexandre and Bakker, Jan, "Noninvasive monitoring of peripheral perfusion," *Intensive care medicine*, vol. 31, no. 10, pp. 1316–1326, 2005.

[84] Backer, Daniel De and Creteur, Jacques, "Regional hypoxia and partial pressure of carbon dioxide gradients: what is the link?," *Intensive care medicine*, vol. 29, no. 12, pp. 2116–2118, 2003.

[85] Tremper, Kevin K and Shoemaker, William C and Shippy, Clay R and Nolan, Lorene S, "Transcutaneous PCO2 monitoring on adult patients in the ICU and the operating room," *Critical care medicine*, vol. 9, no. 10, pp. 752–755, 1981.

[86] Tremper, Kevin K and Mentelos, Richard A and Shoemaker, William C, "Effect of hypercarbia and shock on transcutaneous carbon dioxide at different electrode temperatures," *Critical care medicine*, vol. 8, no. 11, pp. 608–612, 1980.

[87] Eberhard, P and Mindt, W and Schäfer, R, "Cutaneous blood gas monitoring in the adult," *Critical care medicine*, vol. 9, no. 10, pp. 702–705, 1981.

[88] White, Rodney A and Nolan, Lorene and Harley, Daniel and Long, John and Klein, Stanley and Tremper, Kevin and Nelson, Ronald and Tabrisky, Joseph and Shoemaker, William, "Noninvasive evaluation of peripheral vascular disease using transcutaneous oxygen tension," *The American Journal of Surgery*, vol. 144, no. 1, pp. 68–75, 1982.

[89] Gottrup, Finn and Firmin, Richard and Rabkin, John and Halliday, Betty J and Hunt, Thomas K, "Directly measured tissue oxygen tension and arterial oxygen tension assess tissue perfusion," *Critical care medicine*, vol. 15, no. 11, pp. 1030–1036, 1987.

[90] Gisiger, PA and Palma, JP and Eberhard, P, "OxiCarbo®", a single sensor for the non-invasive measurement of arterial oxygen saturation and CO_2partial pressure at the ear lobe," *Sensors and Actuators B: Chemical*, vol. 76, no. 1, pp. 527–530, 2001.

[91] Stücker, M and Struk, A and Altmeyer, P and Herde, M and Baumgärtl, H and Lübbers, DW, "The cutaneous uptake of atmospheric oxygen contributes significantly to the oxygen supply of human dermis and epidermis," *The Journal of physiology*, vol. 538, no. 3, pp. 985–994, 2002.

[92] Eickhoff, JH and Jacobsen, E and others, "Is transcutaneous oxygen tension independent of variations in blood flow and in arterial blood pressure?," *Biotelemetry and patient monitoring*, vol. 9, no. 3, p. 175, 1982.

[93] Peabody, JL and Willis, MM and Gregory, GA and Tooley, WH and Lucey, JF, "Clinical limitations and advantages of transcutaneous oxygen electrodes," *Acta Anaesthesiologica Scandinavica*, vol. 22, no. s68, pp. 76–82, 1978.

[94] Durand, Manuel and Ramanathan, Rangasamy, "Pulse oximetry for continuous oxygen monitoring in sick newborn infants," *The Journal of pediatrics*, vol. 109, no. 6, pp. 1052–1056, 1986.

[95] Newman, Paul G and Rozycki, Grace S, "The history of ultrasound," *Surgical clinics of north America*, vol. 78, no. 2, pp. 179–195, 1998.

[96] Tranquart, F and Grenier, N and Eder, V and Pourcelot, L, "Clinical use of ultrasound tissue harmonic imaging," *Ultrasound in medicine & biology*, vol. 25, no. 6, pp. 889–894, 1999.

[97] Rubin, Jonathan M and Bude, Ronald O and Carson, Paul L and Bree, Robert L and Adler, Ronald S, "Power Doppler US: a potentially useful alternative to mean frequency-based color Doppler US.," *Radiology*, vol. 190, no. 3, pp. 853–856, 1994.

[98] Guiot, C and Gaglioti, P and Oberto, M and Piccoli, E and Rosato, R and Todros, T, "Is three-dimensional power Doppler ultrasound useful in the assessment of placental perfusion in normal and growth-restricted pregnancies?," *Ultrasound in Obstetrics & Gynecology*, vol. 31, no. 2, pp. 171–176, 2008.

[99] Hirsch, W and Bell, DS and Crawford, DC and Kale, SG and McCready, VR and Bamber, JC, "Colour Doppler image analysis for tissue vascularity and perfusion: a preliminary clinical evaluation," *Ultrasound in medicine & biology*, vol. 21, no. 9, pp. 1107–1117, 1995.

[100] Newman, Joel S and Adler, RS and Bude, RO and Rubin, JM, "Detection of soft-tissue hyperemia: value of power Doppler sonography.," *AJR. American journal of roentgenology*, vol. 163, no. 2, pp. 385–389, 1994.

[101] Quaia, Emilio, "Microbubble ultrasound contrast agents: an update," *European radiology*, vol. 17, no. 8, pp. 1995–2008, 2007.

[102] Dymling, Stephan O and Persson, Hans W and Hertz, C Hellmuth, "Measurement of blood perfusion in tissue using Doppler ultrasound," *Ultrasound in medicine & biology*, vol. 17, no. 5, pp. 433–444, 1991.

[103] Eriksson, Roger and Persson, Hans W and Dymling, Stephan O and Lindström, Kjell, "Evaluation of Doppler ultrasound for blood perfusion measurements," *Ultrasound in medicine & biology*, vol. 17, no. 5, pp. 445–452, 1991.

[104] Jansson, T and Persson, HW and Lindström, K, "Estimation of blood perfusion using ultrasound," *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, vol. 213, no. 2, pp. 91–106, 1999.

[105] Wiedeman, Mary Purcell, *An introduction to microcirculation*, vol. 2. Academic Press, 1981.

[106] Johnson, Steven A and Greenleaf, James F and Hansen, CR and Samayoa, WF and Tanaka, M and Lent, A and Christensen, DA and Woolley, RL, "Reconstructing three-dimensional fluid velocity vector fields from acoustic transmission

measurements," in *7th International Symposium on Acoustical Holography and Imaging*, vol. 1, pp. 307–326, 1977.

[107] Marks, Fay A and Tomlinson, Harold W and Brooksby, Glen W, "Comprehensive approach to breast cancer detection using light: photon localization by ultrasound modulation and tissue characterization by spectral discrimination," in *OE/LASE'93: Optics, Electro-Optics, & Laser Applications in Science& Engineering*, pp. 500–510, International Society for Optics and Photonics, 1993.

[108] Wang, Lihong and Jacques, Steven L and Zhao, Xuemei, "Continuous-wave ultrasonic modulation of scattered laser light to image objects in turbid media," *Optics letters*, vol. 20, no. 6, pp. 629–631, 1995.

[109] Leutz, W and Maret, Georg, "Ultrasonic modulation of multiply scattered light," *Physica B: Condensed Matter*, vol. 204, no. 1, pp. 14–19, 1995.

[110] Wang, Lihong V and Ku, Geng, "Frequency-swept ultrasound-modulated optical tomography of scattering media," *Optics letters*, vol. 23, no. 12, pp. 975–977, 1998.

[111] Fang, Hui and Maslov, Konstantin and Wang, Lihong V, "Photoacoustic Doppler effect from flowing small light-absorbing particles," *Physical review letters*, vol. 99, no. 18, p. 184501, 2007.

[112] Fang, Hui and Maslov, Konstantin and Wang, Lihong V, "Photoacoustic Doppler flow measurement in optically scattering media," *Applied Physics Letters*, vol. 91, no. 26, pp. 264103–264103, 2007.

[113] Yao, Junjie and Wang, Lihong V, "Transverse flow imaging based on photoacoustic Doppler bandwidth broadening," *Journal of biomedical optics*, vol. 15, no. 2, pp. 021304–021304, 2010.

[114] Brunker, J and Beard, P, "Pulsed photoacoustic Doppler flowmetry using a cross correlation method," in *BiOS*, pp. 756426–756426, International Society for Optics and Photonics, 2010.

[115] Zhang, Hao F and Maslov, Konstantin and Stoica, George and Wang, Lihong V, "Functional photoacoustic microscopy for high-resolution and noninvasive in vivo imaging," *Nature biotechnology*, vol. 24, no. 7, pp. 848–851, 2006.

[116] Holloway, G Allen and Watkins, Dennis W, "Laser Doppler measurement of cutaneous blood flow," *Journal of Investigative Dermatology*, vol. 69, no. 3, pp. 306–309, 1977.

[117] Nilsson, Gert E and Tenland, Torsten and Oberg, P Ake, "Evaluation of a laser Doppler flowmeter for measurement of tissue blood flow," *Biomedical Engineering, IEEE Transactions on*, no. 10, pp. 597–604, 1980.

[118] Skarphedinsson, JO and Hårding, H and Thoren, P, "Repeated measurements of cerebral blood flow in rats. Comparisons between the hydrogen clearance method and laser Doppler flowmetry," *Acta Physiologica Scandinavica*, vol. 134, no. 1, pp. 133–142, 1988.

[119] Braverman, Irwin M and Keh, Agnes and Goldminz, David, "Correlation of laser Doppler wave patterns with underlying microvascular anatomy," *Journal of investigative dermatology*, vol. 95, no. 3, pp. 283–286, 1990.

[120] Kuznetsova, Larisa V and Tomasek, Nicole and Sigurdsson, Gisli H and Banic, Andrej and Erni, Dominique and Wheatley, Anthony M, "Dissociation between volume blood flow and laser-Doppler signal from rat muscle during changes in vascular tone," *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 274, no. 4, pp. H1248–H1254, 1998.

[121] Kubli, Sandrine and Waeber, Bernard and Dalle-Ave, Anne and Feihl, François, "Reproducibility of laser Doppler imaging of skin blood flow as a tool to assess endothelial function," *Journal of cardiovascular pharmacology*, vol. 36, no. 5, pp. 640–648, 2000.

[122] Bircher, Andreas and Boer, Edith M and Agner, Tove and Wahlberg, Jan E and Serup, Jørgen, "Guidelines for measurement of cutaneous blood flow by laser Doppler flowmetry," *Contact Dermatitis*, vol. 30, no. 2, pp. 65–72, 1994.

[123] Bonner, R and Nossal, R, "Model for laser Doppler measurements of blood flow in tissue," *Applied optics*, vol. 20, no. 12, pp. 2097–2107, 1981.

[124] Maret, G. and Wolf, P. E., "Multiple light scattering from disordered media. The effect of brownian motion of scatterers," *Zeitschrift fur Physik B Condensed Matter*, vol. 65, pp. 409–413, Dec. 1987.

[125] Brown, W., *Dynamic light scattering: the method and some applications*. Monographs on the physics and chemistry of materials, Clarendon Press, 1993.

[126] Boas, DA and Campbell, LE and Yodh, AG, "Scattering and imaging with diffusing temporal field correlations," *Physical Review Letters*, vol. 75, no. 9, p. 1855, 1995.

[127] T. G. Mason and Hu Gang and D. A. Weitz, "Diffusing-wave-spectroscopy measurements of viscoelasticity of complex fluids," *J. Opt. Soc. Am. A*, vol. 14, pp. 139–149, Jan 1997.

[128] Cheung, Cecil and Culver, Joseph P and Takahashi, Kasushi and Greenberg, Joel H and Yodh, AG, "In vivo cerebrovascular measurement combining diffuse near-infrared absorption and correlation spectroscopies," *Physics in medicine and biology*, vol. 46, no. 8, p. 2053, 2001.

[129] Zhou, Chao, *In-vivo optical imaging and spectroscopy of cerebral hemodynamics*. PhD thesis, Penn State University, 2007.

[130] Yu, Guoqiang and Floyd, Thomas F and Durduran, Turgut and Zhou, Chao and Wang, Jiongjiong and Detre, John A and Yodh, Arjun G, "Validation of diffuse correlation spectroscopy for muscle blood flow with concurrent arterial spin labeled perfusion MRI," *Optics express*, vol. 15, no. 3, pp. 1064–1075, 2007.

[131] Mesquita, Rickson C and Durduran, Turgut and Yu, Guoqiang and Buckley, Erin M and Kim, Meeri N and Zhou, Chao and Choe, Regine and Sunar, Ulas and Yodh, Arjun G, "Direct measurement of tissue blood flow and metabolism with diffuse optics," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 369, no. 1955, pp. 4390–4406, 2011.

[132] Chen, Zhongping and Milner, Thomas E and Srinivas, Shyam and Wang, Xiaojun and Malekafzali, Arash and van Gemert, Martin JC and Nelson, J Stuart, "Noninvasive imaging of *in vivo* blood flow velocity using optical Doppler tomography," *Optics letters*, vol. 22, no. 14, pp. 1119–1121, 1997.

[133] Bizheva, Kostadinka K and Siegel, Andy M and Boas, David A, "Path-length-resolved dynamic light scattering in highly scattering random media: The transition to diffusing wave spectroscopy," *Physical Review E*, vol. 58, no. 6, p. 7664, 1998.

[134] Wax, Adam and Yang, Changhuei and Dasari, Ramachandra R and Feld, Michael S, "Path-length-resolved dynamic light scattering: modeling the transition from single to diffusive scattering," *Applied optics*, vol. 40, no. 24, pp. 4222–4227, 2001.

[135] Liu, Gangjun and Jia, Wangcun and Sun, Victor and Choi, Bernard and Chen, Zhongping and others, "High-resolution imaging of microvasculature in human skin in-vivo with optical coherence tomography," *Optics express*, vol. 20, no. 7, p. 7694, 2012.

[136] Choi, Woo June and Na, Jihoon and Eom, Jonghyun and Lee, Byeong Ha, "Stroboscopic wide-field optical coherence gated imaging," in *Proc. of SPIE Vol*, vol. 7169, pp. 71691L–1, 2009.

[137] Zhang, Jun and Chen, Zhongping, "In vivo blood flow imaging by a swept laser source based Fourier domain optical Doppler tomography," *Optics express*, vol. 13, no. 19, pp. 7449–7457, 2005.

[138] Fick, A, "Ueber die Aenderung der Elasticität des Muskels während der Zuckung," *Pflügers Archiv European Journal of Physiology*, vol. 4, no. 1, pp. 301–315, 1871.

[139] Lima, Alexandre and van Bommel, Jasper and Sikorska, Karolina and van Genderen, Michel and Klijn, Eva and Lesaffre, Emmanuel and Ince, Can and Bakker, Jan, "The relation of near-infrared spectroscopy with changes in peripheral circulation in critically ill patients*," *Critical care medicine*, vol. 39, no. 7, pp. 1649–1654, 2011.

[140] Scholkmann, Felix and Kleiser, Stefan and Metz, Andreas and Zimmermann, Raphael and Pavia, Juan Mata and Wolf, Ursula and Wolf, Martin, "A review on continuous wave functional near-infrared spectroscopy and imaging instrumentation and methodology," *NeuroImage*, 2013.

[141] Hueber, Dennis M and Fantini, Sergio and Cerussi, Albert E and Barbieri, Beniamino B, "New optical probe designs for absolute (self-calibrating) NIR tissue hemoglobin measurements," in *BiOS'99 International Biomedical Optics Symposium*, pp. 618–631, International Society for Optics and Photonics, 1999.

[142] Matcher, Steven J and Kirkpatrick, Peter J and Nahid, K and Cope, Mark and Delpy, David T, "Absolute quantification methods in tissue near-infrared spectroscopy," in *Photonics West'95*, pp. 486–495, International Society for Optics and Photonics, 1995.

[143] A. M. S. B.V., "Artinis medical systems website." `www.artinis.com/`, August 2013.

[144] McKinley, Bruce A and Marvin, Robert G and Cocanour, Christine S and Moore, Frederick A, "Tissue hemoglobin O2 saturation during resuscitation of traumatic shock monitored using near infrared spectrometry," *The Journal of Trauma and Acute Care Surgery*, vol. 48, no. 4, pp. 637–642, 2000.

[145] Boushel, R and Langberg, H and Olesen, J and Gonzales-Alonzo, J and Bülow, J and Kjaer, M, "Monitoring tissue oxygen availability with near infrared spec-

troscopy (NIRS) in health and disease," *Scandinavian journal of medicine & science in sports*, vol. 11, no. 4, pp. 213–222, 2001.

[146] Matcher, SJ and Elwell, CE and Cooper, CE and Cope, M and Delpy, DT, "Performance comparison of several published tissue near-infrared spectroscopy algorithms," *Analytical biochemistry*, vol. 227, no. 1, pp. 54–68, 1995.

[147] Duncan, Arlene and Meek, Judith H and Clemence, Matthew and Elwell, Clare E and Tyszczuk, Lidia and Cope, Mark and Delpy, D, "Optical pathlength measurements on adult head, calf and forearm and the head of the newborn infant using phase resolved optical spectroscopy," *Physics in medicine and biology*, vol. 40, no. 2, p. 295, 1995.

[148] Kohl, Matthias and Nolte, Christian and Heekeren, Hauke R and Horst, Susanne and Scholz, Udo and Obrig, Hellmuth and Villringer, Arno, "Determination of the wavelength dependence of the differential pathlength factor from near-infrared pulse signals," *Physics in Medicine and Biology*, vol. 43, no. 6, p. 1771, 1998.

[149] Essenpreis, Matthias and Elwell, CE and Cope, M and Van der Zee, P and Arridge, SR and Delpy, DT, "Spectral dependence of temporal point spread functions in human tissues," *Applied Optics*, vol. 32, no. 4, pp. 418–425, 1993.

[150] H. T. Inc., "Hutchinson technology website website." `http://www.htibiomeasurement.com/`, August 2013.

[151] James Housley, *Determination of the optimum metrics for the characterisation of scattering media.* PhD thesis, University of Nottingham, July 2010.

[152] R. J. Cooper, E. Magee, N. Everdell, S. Magazov, M. Varela, D. Airantzis, A. P. Gibson, and J. C. Hebden, "Monstir ii: A 32-channel, multispectral, time-resolved optical tomography system for neonatal brain imaging," *Review of Scientific Instruments*, vol. 85, no. 5, p. 053105, 2014.

[153] S. Donati, "Photodetectors: Devices, circuits, and applications," *Measurement Science and Technology*, vol. 12, no. 5, p. 653, 2001.

[154] L. D. Sieno, D. Contini, A. D. Mora, A. Torricelli, L. Spinelli, R. Cubeddu, A. Tosi, G. Boso, and A. Pifferi, "Functional near-infrared spectroscopy at small source-detector distance by means of high dynamic-range fast-gated spad acquisitions: First in-vivo measurements," in *Neurophotonics*, p. 880402, Optical Society of America, 2013.

[155] M. Ghioni, A. Gulinatti, I. Rech, P. Maccagnani, and S. Cova, "Large-area low-jitter silicon single photon avalanche diodes," 2008.

[156] A. Stewart, V. Saveliev, S. Bellis, D. Herbert, P. Hughes, and J. Jackson, "Performance of 1-mm 2 silicon photomultiplier," *Quantum Electronics, IEEE Journal of*, vol. 44, no. 2, pp. 157–164, 2008.

[157] "ST microelectronics website." `http://www.st.com/web/catalog/mmc/FM132/CL2136/SC1934/PF260441`, December 2014.

[158] Tualle, J-M and Nghiêm, Ha Lien and Cheikh, Monia and Ettori, Dominique and Tinet, Eric and Avrillier, Sigrid, "Time-resolved diffusing wave spectroscopy beyond 300 transport mean free paths," *JOSA A*, vol. 23, no. 6, pp. 1452–1457, 2006.

[159] Tualle, J-M and Tinet, E and Avrillier, S, "A new and easy way to perform time-resolved measurements of the light scattered by a turbid medium," *Optics communications*, vol. 189, no. 4, pp. 211–220, 2001.

[160] Tualle, J-M and Nghiem, Ha Lien and Schäfauer, Christof and Berthaud, Patrice and Tinet, Eric and Ettori, Dominique and Avrillier, Sigrid, "Time-resolved measurements from speckle interferometry," *Optics letters*, vol. 30, no. 1, pp. 50–52, 2005.

[161] Tualle, Jean-Michel and others, "Method for analyzing a diffusing sample by time resolution measurement," June 7 2005. US Patent 6,903,825.

[162] "Amateur laser holography website." `http://http://redlum.xohp.pagesperso-orange.fr/laser/modeanalysis.html`, August 2013.

[163] Hawthorn, CJ and Weber, KP and Scholten, RE, "Littrow configuration tunable external cavity diode laser with fixed direction output beam," *Review of scientific instruments*, vol. 72, no. 12, pp. 4477–4479, 2001.

[164] Van Beekvelt, Mireille CP and Colier, Willy NJM and Wevers, Ron A and Van Engelen, Baziel GM, "Performance of near-infrared spectroscopy in measuring local O2 consumption and blood flow in skeletal muscle," *Journal of Applied Physiology*, vol. 90, no. 2, pp. 511–519, 2001.

[165] De Blasi, Roberto A and Ferrari, Marco and Natali, Andrea and Conti, Giorgio and Mega, Annamaria and Gasparetto, Alessandro, "Noninvasive measurement of forearm blood flow and oxygen consumption by near-infrared spectroscopy," *Journal of Applied Physiology*, vol. 76, no. 3, pp. 1388–1393, 1994.

[166] Raynaud, JS and Duteil, S and Vaughan, JT and Hennel, F and Wary, C and Leroy-Willig, A and Carlier, PG, "Determination of skeletal muscle perfusion using arterial spin labeling NMRI: validation by comparison with venous occlusion plethysmography," *Magnetic resonance in medicine*, vol. 46, no. 2, pp. 305–311, 2001.

[167] Woinarski, Nicholas CZ and Suzuki, Satoshi and Lipcsey, Miklos and Lumsden, Natalie and Chin-Dusting, Jaye and Schneider, Antoine G and Bailey, Michael and Bellomo, Rinaldo, "Near-infrared spectroscopy of the thenar eminence to estimate forearm blood flow," *Critical Care and Resuscitation*, vol. 15, no. 4, p. 323, 2013.

[168] Bollinger, A and Butti, P and Barras, J-P and Trachsler, H and Siegenthaler, W, "Red blood cell velocity in nailfold capillaries of man measured by a television microscopy technique," *Microvascular research*, vol. 7, no. 1, pp. 61–72, 1974.

[169] Dobbe, Johannes GG and Streekstra, Geert J and Atasever, Bektas and Van Zijderveld, R and Ince, Can, "Measurement of functional microcirculatory geometry and velocity distributions using automated image analysis," *Medical & biological engineering & computing*, vol. 46, no. 7, pp. 659–670, 2008.

[170] Groner, Warren and Winkelman, James W and Harris, Anthony G and Ince, Can and Bouma, Gerrit J and Messmer, Konrad and Nadeau, Richard G, "Orthogonal polarization spectral imaging: a new method for study of the microcirculation," *Nature medicine*, vol. 5, no. 10, pp. 1209–1212, 1999.

[171] Mathura, Keshen R and Vollebregt, Karlijn C and Boer, Kees and De Graaff, Jurgen C and Ubbink, Dirk T and Ince, Can, "Comparison of OPS imaging and conventional capillary microscopy to study the human microcirculation," *Journal of Applied Physiology*, vol. 91, no. 1, pp. 74–78, 2001.

[172] Goedhart, PT and Khalilzada, M and Bezemer, R and Merza, J and Ince, C, "Sidestream Dark Field (SDF) imaging: a novel stroboscopic LED ring-based imaging modality for clinical assessment of the microcirculation," *Optics express*, vol. 15, no. 23, pp. 15101–15114, 2007.

[173] Welzel, Julia and Lankenau, Eva and Birngruber, Reginald and Engelhardt, Ralf, "Optical coherence tomography of the human skin," *Journal of the American Academy of Dermatology*, vol. 37, no. 6, pp. 958–963, 1997.

[174] White, Brian and Pierce, Mark and Nassif, Nader and Cense, Barry and Park, B and Tearney, Guillermo and Bouma, Brett and Chen, Teresa and de Boer, Johannes, "In vivo dynamic human retinal blood flow imaging using ultra-high-

speed spectral domain optical coherence tomography," *Optics Express*, vol. 11, no. 25, pp. 3490–3497, 2003.

[175] Barton, Jennifer and Stromski, Steven, "Flow measurement without phase information in optical coherence tomography images," *Opt. Express*, vol. 13, no. 14, pp. 5234–5239, 2005.

[176] Jonathan, Enock and Enfield, Joey and Leahy, Martin J, "Correlation mapping method for generating microcirculation morphology from optical coherence tomography (OCT) intensity images," *Journal of biophotonics*, vol. 4, no. 9, pp. 583–587, 2011.

[177] Doronin, Alexander and Meglinski, Igor, "Imaging of subcutaneous microcirculation vascular network by double correlation Optical Coherence Tomography," *Laser & Photonics Reviews*, 2013.

[178] Yousefi, Siavash and Qin, Jia and Wang, Ruikang K, "Super-resolution spectral estimation of optical micro-angiography for quantifying blood flow within microcirculatory tissue beds in vivo," *Biomedical optics express*, vol. 4, no. 7, p. 1214, 2013.

[179] Felice, Claudio and Latini, Giuseppe and Vacca, Paola and Kopotic, Robert J, "The pulse oximeter perfusion index as a predictor for high illness severity in neonates," *European journal of pediatrics*, vol. 161, no. 10, pp. 561–562, 2002.

[180] Zaramella, Patrizia and Freato, Federica and Quaresima, Valentina and Ferrari, Marco and Vianello, Andrea and Giongo, Diego and Conte, Lorena and Chiandetti, Lino, "Foot pulse oximeter perfusion index correlates with calf muscle perfusion measured by near-infrared spectroscopy in healthy neonates," *Journal of Perinatology*, vol. 25, no. 6, pp. 417–422, 2005.

[181] Diab, Mohamed K and Kiani-Azarbayjany, Esmaiel, "Signal processing apparatus and method," Jan. 9 1996. US Patent 5,482,036.

[182] Goldman, Julian and Petterson, Michael and Kopotic, Robert and Barker, Steven, "Masimo Signal Extraction Pulse Oximetry," *Journal of Clinical Monitoring and Computing*, vol. 16, pp. 475–483, 2000. 10.1023/A:1011493521730.

[183] Chan, Fang Chiat Daniel and Hayes, Matthew James and Smith, Peter Richard and others, "Venous pulse oximetry," Aug. 28 2007. US Patent 7,263,395.

[184] Nilsson, Lena and Johansson, Anders and Kalman, Sigga, "Respiratory variations in the reflection mode photoplethysmographic signal. Relationships to peripheral

venous pressure," *Medical and Biological Engineering and Computing*, vol. 41, no. 3, pp. 249–254, 2003.

[185] Webster, John, *Design of pulse oximeters*. Bristol Philadelphia: Institute of Physics Pub, 1997.

[186] Patterson, Michael S and Chance, Britton and Wilson, Brian C and others, "Time resolved reflectance and transmittance for the non-invasive measurement of tissue optical properties," *Appl. Opt*, vol. 28, no. 12, pp. 2331–2336, 1989.

[187] Housley, James, *Determination of the optimum metrics for the characterisation of scattering media*. PhD thesis, University of Nottingham, 2010.

[188] Elaloufi, Rachid and Carminati, Rémi and Greffet, Jean-Jacques, "Definition of the diffusion coefficient in scattering and absorbing media," *JOSA A*, vol. 20, no. 4, pp. 678–685, 2003.

[189] Flock, Stephen T and Patterson, Michael S and Wilson, Brian C and Wyman, Douglas R, "Monte Carlo modeling of light propagation in highly scattering tissues. I. Model predictions and comparison with diffusion theory," *Biomedical Engineering, IEEE Transactions on*, vol. 36, no. 12, pp. 1162–1168, 1989.

[190] Bashkatov, AN and Genina, EA and Kochubey, VI and Tuchin, VV, "Optical properties of human skin, subcutaneous and mucous tissues in the wavelength range from 400 to 2000 nm," *Journal of Physics D: Applied Physics*, vol. 38, no. 15, p. 2543, 2005.

[191] Crowe, John A and Damianou, Damianos, "The Wavelength Dependence of the Photoplethysmogram and its implication to Pulse Oximetry," in *Engineering in Medicine and Biology Society, 1992 14th Annual International Conference of the IEEE*, vol. 6, pp. 2423–2424, IEEE, 1992.

[192] Damianou, Damianos, *The wavelength dependence of the photoplethysmogram and its implication to pulse oximetry.* PhD thesis, University of Nottingham, 1995.

[193] Lobos, Anna-Theresa and Menon, Kusum, "A multidisciplinary survey on capillary refill time: Inconsistent performance and interpretation of a common clinical test," *Pediatric Critical Care Medicine*, vol. 9, no. 4, pp. 386–391, 2008.

[194] Brown, Lawrence H and Prasad, N Heramba and Whitley, Theodore W, "Adverse lighting condition effects on the assessment of capillary refill," *The American journal of emergency medicine*, vol. 12, no. 1, pp. 46–47, 1994.

[195] Messerges, Joanne L and Hutchinson, George M, "Method and apparatus for measuring capillary refill time and carrying out pulse oximetry with a single device," Dec. 20 2011. US Patent 8,082,017.

[196] Kviesis-Kipge, E and Curkste, E and Spigulis, J and Gardovska, D, "Optical studies of the capillary refill kinetics in fingertips," in *World Congress on Medical Physics and Biomedical Engineering, September 7-12, 2009, Munich, Germany*, pp. 377–379, Springer, 2009.

[197] Shavit, Itai and Brant, Rollin and Nijssen-Jordan, Cheri and Galbraith, Roger and Johnson, David W, "A novel imaging technique to measure capillary-refill time: improving diagnostic accuracy for dehydration in young children with gastroenteritis," *Pediatrics*, vol. 118, no. 6, pp. 2402–2408, 2006.

[198] Takayesu, James Kimo and Lozner, Alison Wiley, "Pediatrics, Dehydration," *Pediatrics*, vol. 8, p. 18, 1910.

[199] Guedel, Arthur E, "Cyclopropane anesthesia," *Anesthesiology*, vol. 1, no. 1, pp. 13–25, 1940.

[200] Crismon, JM and Fuhrman, FA, "Studies on gangrene following cold injury: vi. Capillary blood flow after cold injury, the effects of rapid warming, and sympathetic block," *Journal of Clinical Investigation*, vol. 26, no. 3, p. 468, 1947.

[201] Champion, Howard R and Sacco, William J and Carnazzo, Anthony J and Copes, Wayne and Fouty, William J, "Trauma score," *Critical Care Medicine*, vol. 9, no. 9, pp. 672–676, 1981.

[202] Baraff, Larry J, "Capillary refill: is it a useful clinical sign?," *Pediatrics*, vol. 92, no. 5, pp. 723–724, 1993.

[203] Schriger, David L and Baraff, Larry, "Defining normal capillary refill: variation with age, sex, and temperature," *Annals of emergency medicine*, vol. 17, no. 9, pp. 932–935, 1988.

[204] Gorelick, Marc H and Shaw, Kathy N and Baker, M Douglas, "Effect of ambient temperature on capillary refill in healthy children," *Pediatrics*, vol. 92, no. 5, pp. 699–702, 1993.

[205] Champion, Howard R and Sacco, William J and Copes, Wayne S and Gann, Donald S and Gennarelli, Thomas A and Flanagan, Maureen E, "A revision of the Trauma Score," *The Journal of Trauma and Acute Care Surgery*, vol. 29, no. 5, pp. 623–629, 1989.

[206] Saavedra, Jose M and Harris, Glenn D and Li, Song and Finberg, Laurence, "Capillary refilling (skin turgor) in the assessment of dehydration," *Archives of Pediatrics & Adolescent Medicine*, vol. 145, no. 3, p. 296, 1991.

[207] Gray, Katherine and Briseno, Michael R and Otsuka, Norman Y, "The association between capillary refill time and arterial flow in the pediatric upper extremity," *Journal of Pediatric Orthopaedics B*, vol. 17, no. 5, pp. 257–260, 2008.

[208] Pickard, Amelia and Karlen, Walter and Ansermino, J Mark, "Capillary refill time: is it still a useful clinical sign?," *Anesthesia & Analgesia*, vol. 113, no. 1, pp. 120–123, 2011.

[209] Lima, Alexandre and van Genderen, Michel Egide and Klijn, Eva and Bakker, Jan and van Bommel, Jasper, "Peripheral vasoconstriction influences thenar oxygen saturation as measured by near-infrared spectroscopy," *Intensive care medicine*, vol. 38, no. 4, pp. 606–611, 2012.

[210] Saaty, Thomas L, "Optimization by the Analytic Hierarchy Process," tech. rep., Air Force Office of Scientific Research, Bolling AFB DC, 1979.

[211] Saaty, Thomas L and Kearns, Kevin P, *Analytical planning: the organization of systems*. RWS Publications, 1985.

[212] Saaty, Thomas L, "How to make a decision: the analytic hierarchy process," *European journal of operational research*, vol. 48, no. 1, pp. 9–26, 1990.

[213] Morgan, Jim, "Manufacturing Readiness Levels (MRLs) and Manufacturing Readiness Assessments (MRAs)," tech. rep., DTIC Document, 2008.

[214] Siegel, A and Marota, JJ and Boas, David, "Design and evaluation of a continuous-wave diffuse optical tomography system," *Optics Express*, vol. 4, no. 8, pp. 287–298, 1999.

[215] Farrell, Thomas J and Patterson, Michael S and Wilson, Brian, "A diffusion theory model of spatially resolved, steady-state diffuse reflectance for the noninvasive determination of tissue optical properties in vivo," *Medical physics*, vol. 19, p. 879, 1992.

[216] Amaboldi, C and Bucci, C and Cremonesi, O and Fascilla, A and Nucciotti, A and Pavan, M and Pessina, G and Pirro, S and Previtali, E and Sisti, M, "Low-frequency noise characterization of very large value resistors," *Nuclear Science, IEEE Transactions on*, vol. 49, no. 4, pp. 1808–1813, 2002.

[217] Bolin, Frank P and Preuss, Luther E and Taylor, Roy C and Ference, Robert J, "Refractive index of some mammalian tissues using a fiber optic cladding method," *Applied Optics*, vol. 28, no. 12, pp. 2297–2303, 1989.

[218] Tuchin, Valery V and Culver, Joseph P and Cheung, Cecil and Tatarkova, Svetlana A and DellaVecchia, Michael A and Zimnyakov, Dmitry A and Chaussky, Anatoly A and Yodh, Arjun G and Chance, Britton, "Refractive index matching of tissue components as a new technology for correlation and diffusing-photon spectroscopy and imaging," in *BiOS'99 International Biomedical Optics Symposium*, pp. 111–120, International Society for Optics and Photonics, 1999.

[219] Ziegler, JG and Nichols, NB, "Optimum settings for automatic controllers," *trans. ASME*, vol. 64, no. 11, 1942.

[220] Hendriks, FM and Brokken, Dirk and Van Eemeren, JTWM and Oomens, CWJ and Baaijens, FPT and Horsten, JBAM, "A numerical-experimental method to characterize the non-linear mechanical behaviour of human skin," *Skin research and technology*, vol. 9, no. 3, pp. 274–283, 2003.

[221] Sedgewick, Robert, "Implementing quicksort programs," *Communications of the ACM*, vol. 21, no. 10, pp. 847–857, 1978.

[222] Roberts, David H and Lehár, Joseph and Dreher, John W, "Time series analysis with clean-Part One-Derivation of a spectrum," *The Astronomical Journal*, vol. 93, p. 968, 1987.

[223] Baisch, Stefan and Bokelmann, Götz HR, "Spectral analysis with incomplete time series: an example from seismology," *Computers & Geosciences*, vol. 25, no. 7, pp. 739–750, 1999.

[224] Pradhl, S., "Optical absorption of hemoglobin." `http://omlc.ogi.edu/spectra/hemoglobin/`, December 1999. Oregon Medical Laser Center, Portland, Oreg.

[225] Hyttel-Sorensen, Simon and Sorensen, Line C and Riera, Joan and Greisen, Gorm, "Tissue oximetry: a comparison of mean values of regional tissue saturation, reproducibility and dynamic range of four NIRS-instruments on the human forearm," *Biomedical optics express*, vol. 2, no. 11, p. 3047, 2011.

[226] Brix, Gunnar and Bahner, Malte L and Hoffmann, Ulf and Horvath, Andrea and Schreiber, Wolfgang, "Regional blood flow, capillary permeability, and compartmental volumes: measurement with dynamic CT—initial experience," *Radiology*, vol. 210, no. 1, pp. 269–276, 1999.

[227] Rubins, U and Upmalis, V and Rubenis, O and Jakovels, D and Spigulis, J, "Real-time photoplethysmography imaging system," in *15th Nordic-Baltic Conference on Biomedical Engineering and Medical Physics (NBC 2011)*, pp. 183–186, Springer, 2011.

[228] Chance, Britton, "Spectrophotometric examination of tissue of small dimension," Apr. 4 1995. US Patent 5,402,778.

[229] Boas, David A and Franceschini, Maria Angela and Dunn, Andy K and Strangman, Gary, "Noninvasive imaging of cerebral activation with diffuse optical tomography," *In vivo optical imaging of brain function*, pp. 193–221, 2002.

[230] Chance, Britton and Leigh, JS, "Oxygen intermediates and mixed valence states of cytochrome oxidase: infrared absorption difference spectra of compounds A, B, and C of cytochrome oxidase and oxygen," *Proceedings of the National Academy of Sciences*, vol. 74, no. 11, pp. 4777–4780, 1977.

[231] Matcher, SJ and Cope, M and Delpy, DT, "Use of the water absorption spectrum to quantify tissue chromophore concentration changes in near-infrared spectroscopy," *Physics in medicine and biology*, vol. 39, no. 1, p. 177, 1994.

[232] Van der Zee, P and Cope, M and Arridge, SR and Essenpreis, M and Potter, LA and Edwards, AD and Wyatt, JS and McCormick, DC and Roth, SC and Reynolds, EOR and others, "Experimentally measured optical pathlengths for the adult head, calf and forearm and the head of the newborn infant as a function of inter optode spacing," in *Oxygen Transport to Tissue XIII*, pp. 143–153, Springer, 1992.

[233] van der Sluijs, Marco C and Colier, Willy N and Houston, Ralph JF and Oeseburg, Berend, "New and highly sensitive continuous-wave near-infrared spectrophotometer with multiple detectors," in *BiOS Europe'97*, pp. 63–72, International Society for Optics and Photonics, 1998.

[234] Christoph Zysset and Nassim Nasseri and Lars Büthe and Niko Münzenrieder and Thomas Kinkeldei and Luisa Petti and Stefan Kleiser and Giovanni A. Salvatore and Martin Wolf and Gerhard Tröster, "Textile integrated sensors and actuators for near-infrared spectroscopy," *Opt. Express*, vol. 21, pp. 3213–3224, Feb 2013.

[235] James E. Mark, ed., *Physical Properties of Polymers Handbook*. Springer, 2 ed., 12 2006.

[236] Webster, I, "Recent developments in pressure-sensitive adhesives for medical applications," *International journal of adhesion and adhesives*, vol. 17, no. 1, pp. 69–73, 1997.

[237] "The University of Nottingham code of research conduct and research ethics." `http://www.nottingham.ac.uk/fabs/rgs/documents/code-of-research-conduct-and-research-ethics-approved-january-2010.pdf`, July 2014.

[238] Watson, Adam and Kelly, Anne, "Measuring capillary refill time is useless," *Emergency Medicine*, vol. 5, no. 2, pp. 90–93, 1993.

[239] Abramson, David I and Kahn, Alan and Tuck Jr, Samuel and Turman, George A and Rejal, Habib and Fleischer, Clara J, "Relationship between a range of tissue temperature and local oxygen uptake in the human forearm. I. Changes observed under resting conditions," *Journal of Clinical Investigation*, vol. 37, no. 7, p. 1031, 1958.

[240] Gallen, IW and Macdonald, IA, "Effect of two methods of hand heating on body temperature, forearm blood flow, and deep venous oxygen saturation," *American Journal of Physiology-Endocrinology And Metabolism*, vol. 259, no. 5, pp. E639–E643, 1990.

[241] Detry, JM and Brengelmann, George L and Rowell, Loring B and Wyss, Craig, "Skin and muscle components of forearm blood flow in directly heated resting man.," *Journal of applied physiology*, vol. 32, no. 4, pp. 506–511, 1972.

[242] Peng, C-K and Buldyrev, Sergey V and Havlin, Shlomo and Simons, M and Stanley, H Eugene and Goldberger, Ary L, "Mosaic organization of DNA nucleotides," *Physical Review E*, vol. 49, no. 2, p. 1685, 1994.

[243] Gu, Gao-Feng and Zhou, Wei-Xing, "Detrended fluctuation analysis for fractals and multifractals in higher dimensions," *Physical Review E*, vol. 74, no. 6, p. 061104, 2006.

[244] Malpas, Simon C, "Neural influences on cardiovascular variability: possibilities and pitfalls," *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 282, no. 1, pp. H6–H20, 2002.

[245] White, Halbert, "Using least squares to approximate unknown regression functions," *International Economic Review*, pp. 149–170, 1980.

[246] Cochrane, Donald and Orcutt, Guy H, "Application of least squares regression to relationships containing auto-correlated error terms," *Journal of the American Statistical Association*, vol. 44, no. 245, pp. 32–61, 1949.

[247] Prais, Sigbert J and Winsten, Christopher B, "Trend estimators and serial correlation," tech. rep., Cowles Commission discussion paper, 1954.

[248] Newey, Whitney K and West, Kenneth D, "A simple, positive semi-definite, heteroskedasticity and autocorrelationconsistent covariance matrix," 1986.

[249] Zurlo, Francesco and Larson, Karen and Bogardus, Clifton and Ravussin, Eric, "Skeletal muscle metabolism is a major determinant of resting energy expenditure.," *Journal of Clinical Investigation*, vol. 86, no. 5, p. 1423, 1990.

[250] Cochran, William G, "The distribution of quadratic forms in a normal system, with applications to the analysis of covariance," in *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 30, pp. 178–191, Cambridge Univ Press, 1934.

[251] Lehmann, Erich Leo and Casella, George, *Theory of point estimation*, vol. 31. Springer, 1998.

[252] O'Donnell, CPF and Kamlin, COF and Davis, PG and Morley, CJ, "Obtaining pulse oximetry data in neonates: a randomised crossover study of sensor application techniques," *Archives of Disease in Childhood-Fetal and Neonatal Edition*, vol. 90, no. 1, pp. F84–F85, 2005.

[253] Imadojemu, Virginia A and Gleeson, Kevin and Gray, Kristen S and Sinoway, Lawrence I and Leuenberger, Urs A, "Obstructive apnea during sleep is associated with peripheral vasoconstriction," *American journal of respiratory and critical care medicine*, vol. 165, no. 1, pp. 61–66, 2002.

[254] Garpestad, Erik and Katayama, Hisashi and Parker, J Anthony and Ringler, Jack and Lilly, Jennifer and Yasuda, Tsunehiro and Moore, Richard H and Strauss, H William and Weiss, J Woodrow, "Stroke volume and cardiac output decrease at termination of obstructive apneas," *Journal of Applied Physiology*, vol. 73, no. 5, pp. 1743–1748, 1992.

[255] Johnson, John M, "The cutaneous circulation," in *Laser-Doppler blood flowmetry*, pp. 121–139, Springer, 1990.

[256] Bernardi, Luciano and Hayoz, Daniel and Wenzel, René and Passino, Claudio and Calciati, Alessandro and Weber, Roger and Noll, Georg, "Synchronous and

baroceptor-sensitive oscillations in skin microcirculation: evidence for central autonomic control," *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 42, no. 4, p. H1867, 1997.

[257] Malliani, Alberto and Pagani, Massimo and Lombardi, Federico and Cerutti, Sergio, "Cardiovascular neural regulation explored in the frequency domain.," *Circulation*, vol. 84, no. 2, pp. 482–492, 1991.

[258] Mayer, Sigmund, "Die peripherische Nervenzelle und das sympathische Nervensystem," *European Archives of Psychiatry and Clinical Neuroscience*, vol. 6, no. 2, pp. 353–446, 1876.

[259] Elstad, Maja and Walløe, Lars and Chon, Ki H and Toska, Karin, "Low-frequency fluctuations in heart rate, cardiac output and mean arterial pressure in humans: what are the physiological relationships?," *Journal of hypertension*, vol. 29, no. 7, pp. 1327–1336, 2011.

[260] Zelis, Robert and Longhurst, John and Capone, Robert J and Mason, Dean T and Kleckner, Robert, "A comparison of regional blood flow and oxygen utilization during dynamic forearm exercise in normal subjects and patients with congestive heart failure," *Circulation*, vol. 50, no. 1, pp. 137–143, 1974.

[261] Zelis, Robert and Mason, Dean T and Braunwald, Eugene, "Partition of blood flow to the cutaneous and muscular beds of the forearm at rest and during leg exercise in normal subjects and in patients with heart failure," *Circulation Research*, vol. 24, no. 6, pp. 799–806, 1969.

[262] Scargle, Jeffrey D, "Studies in astronomical time series analysis. II-Statistical aspects of spectral analysis of unevenly spaced data," *The Astrophysical Journal*, vol. 263, pp. 835–853, 1982.

[263] Granger, Clive WJ, "Some recent development in a concept of causality," *Journal of econometrics*, vol. 39, no. 1, pp. 199–211, 1988.

[264] Shore, Angela C, "Capillaroscopy and the measurement of capillary pressure," *British journal of clinical pharmacology*, vol. 50, no. 6, pp. 501–513, 2000.

[265] Levick, JR and Michel, CC, "The effects of position and skin temperature on the capillary pressures in the fingers and toes," *The Journal of physiology*, vol. 274, no. 1, pp. 97–109, 1978.

[266] Braverman, Irwin M, "The cutaneous microcirculation," in *Journal of Investigative Dermatology Symposium Proceedings*, vol. 5, pp. 3–9, Nature Publishing Group, 2000.

[267] Kassab, Ghassan S, "Scaling laws of vascular trees: of form and function," *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 290, no. 2, pp. H894–H903, 2006.

[268] Lund, N and Damon, DH and Damon, DN and Duling, BR, "Capillary grouping in hamster tibials anterior muscles: flow patterns, and physiological significance.," *International journal of microcirculation, clinical and experimental/sponsored by the European Society for Microcirculation*, vol. 5, no. 4, pp. 359–372, 1986.

[269] Berg, Bradley R and Sarelius, Ingrid H, "Functional capillary organization in striated muscle," *American Journal of Physiology*, vol. 268, pp. H1215–H1215, 1995.

[270] Mankins, John C, "Technology readiness levels," 1995.

[271] Harrison, DK and Kessler, M, "Local hydrogen clearance as a method for the measurement of capillary blood flow," *Physics in Medicine and Biology*, vol. 34, no. 10, p. 1413, 1989.

Appendix A

Further details of capillary refill study hardware

A.1 Pneumatic capillary refill sensor used in the PaedCEW study

This sensor was designed by Dr David Morris to be manufactured in significant numbers within the university of Nottingham. For this reason the size was increased considerably to 40 by 45 by 11mm, and a foldered, off the shelf size zero blood pressure cuff was used as the air bladder to apply force to a rigid FR4 PCB. A polyurethane membrane (EU35, Smith and Nephew extruded films Ltd.) was glued across the bottom face of the sensor.

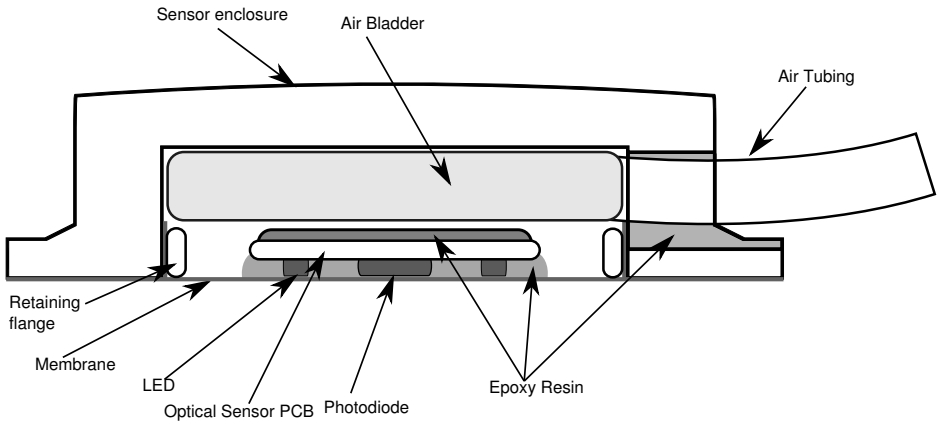


Figure A.1: The capillary refill sensor design used for the PaedCEW study. Image © David Morris, University of Nottingham.

A.2 PaedCEW experimental setup

Figure A.2 shows the experimental setup during a typical baseline recording. The camera based sensor data found to suffer from optical coupling and reflection problems, and was not used as part of the Chapter 10 analysis.

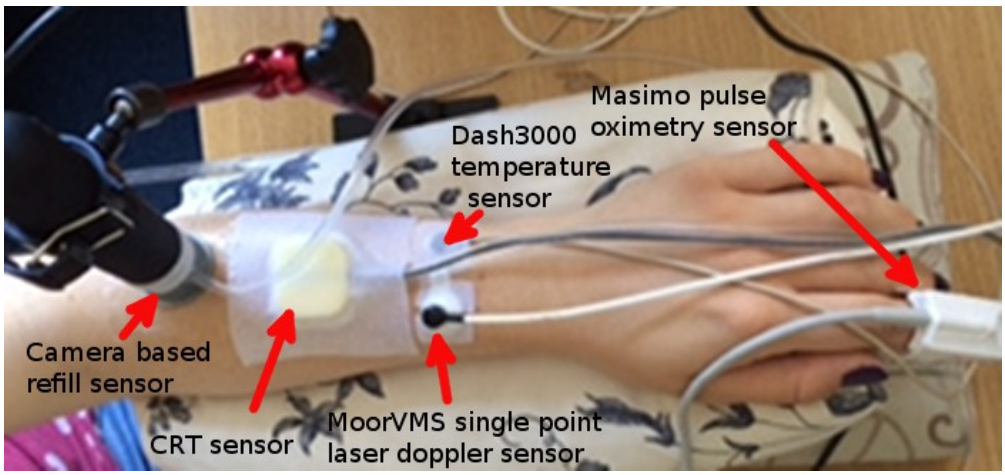


Figure A.2: Example healthy adult volunteer baseline measurement setup for the PaedCEW study. Refill sensor is the same design as above.

Appendix B

Further details of polyurethane capillary refill sensor

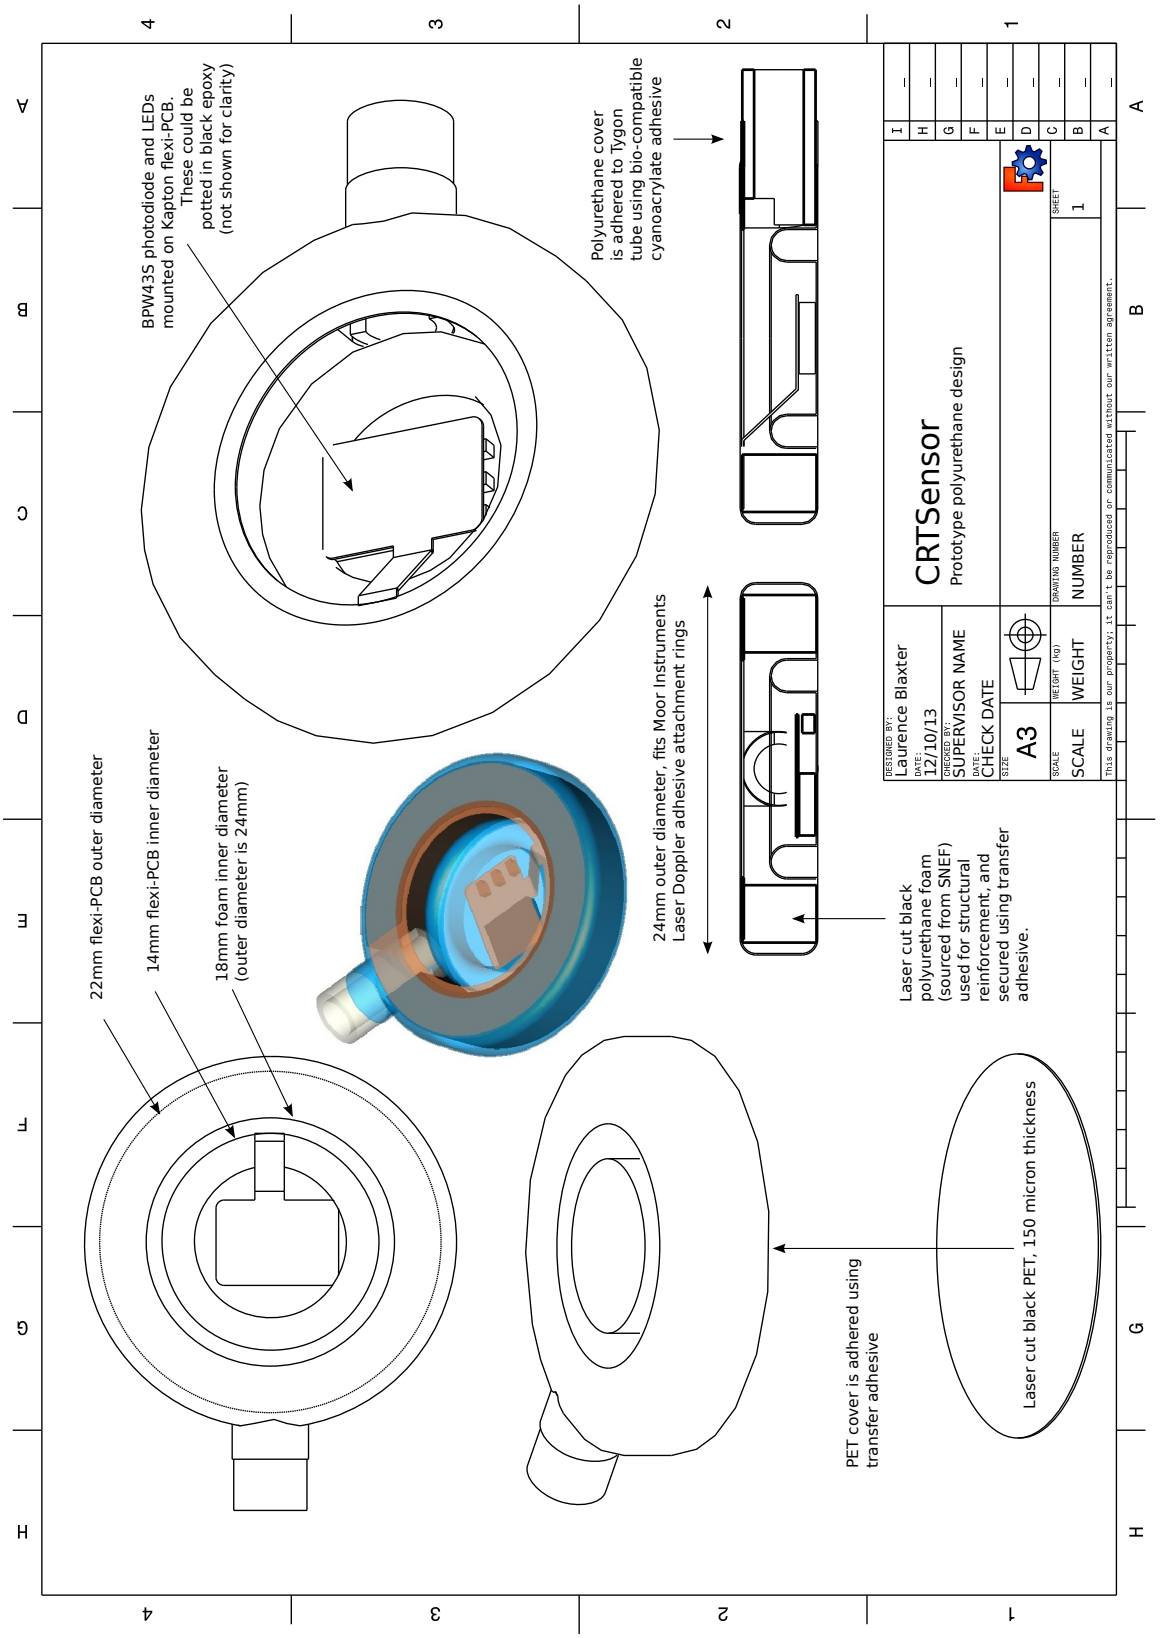


Figure B.1: Thermoformed polyurethane sensor design.